arXiv:math/0611006v2 [math.GR] 2 Dec 2007

Coarse decompositions for boundaries of CAT(0) groups

Dan P. Guralnik
Vanderbilt University

February 2, 2008

Abstract

In this work we introduce a new combinatorial notion of boundary $\Re C$ of an ω-dimensional cubing C. $\Re C$ is defined to be the set of almost-equality classes of ultrafilters on the standard system of halfspaces of C, endowed with an order relation reflecting the interaction between the Tychonoff closures of the classes.

When C arises as the dual of a cubulation – or discrete system of halfspaces – $\mathcal{H}$ of a CAT(0) space X (for example, the Niblo-Reeves cubulation of the Davis-Moussong complex of a finite rank Coxeter group), we show how $\mathcal{H}$ induces a function $\rho : \partial_\infty X \to \Re C$. We develop a notion of uniformness for $\mathcal{H}$, generalizing the parallel walls property enjoyed by Coxeter groups, and show that, if the pair $(X, \mathcal{H})$ admits a geometric action by a group G, then the fibers of ρ form a stratification of $\partial_\infty X$ graded by the order structure of $\Re C$. We also show how this structure computes the components of the Tits boundary of X.

Finally, using our result from another paper, that the uniformness of a cubulation as above implies the local finiteness of C, we give a condition for the co-compactness of the action of G on C in terms of ρ, generalizing a result of Williams, previously known only for Coxeter groups.

1 Introduction

The current research project began with the aim of exploring the relation between the asymptotic topology of a non-positively curved group G and its ability to split (as an amalgam or an HNN extension) over a finitely generated quasi-convex subgroup. A group G is said to be non-positively curved, if it acts properly and co-compactly by isometries (i.e., – *geometrically*) on a CAT(0) space X. The class of CAT(0) groups may be regarded (though the level to which this is true remains an important open question in geometric group theory) as a generalization of the class of (strongly) relatively-hyperbolic groups. For this reason, of particular interest are splittings of CAT(0) groups over finitely generated virtually-abelian subgroups: a hyperbolic group cannot contain a free abelian subgroup of rank 2, while a relatively-hyperbolic group may only contain such a subgroup within a parabolic subgroup. This fact has a serious impact on the connectivity properties of the boundaries of such groups (see [Bow98, Gur05]). In this respect CAT(0)

groups are more flexible, in the sense that they admit abelian subgroups of rank greater than one, though in a controlled fashion: any subgroup $H \cong \mathbb{Z}^d$ in a group G acting geometrically on a CAT(0) space X necessarily stabilizes an isometrically-embedded copy F of $\mathbb{E}^d$, on which it acts by translations so that F/H is isometric to a d-torus.
Another reason for considering finitely generated abelian subgroups as candidates for splittings lies in the important works of Rips and Sela, where it is shown that every finitely presented torsion-free group (initially, every torsion-free word-hyperbolic group) has a canonical decomposition as a graph of groups with virtually abelian edge groups (the "JSJ splitting" of the group), which is, in some sense, maximal among all such decompositions. This fact lies at the base of Sela's approach to constructing algebraic geometry over groups.

Let G denote a finitely generated group and let Γ be a Cayley graph of G with respect to a fixed finite generating set. Note that the natural action of G on Γ by left translations is a geometric action. Suppose now X is a geodesic metric space admitting a geometric action by G; it is then basic to the approach of geometric group theory to identify Γ with X, because, by the Švarc-Milnor lemma, Γ is quasi-isometric to X. As a result, many properties of X, including topological properties of compactifications of X, influence the asymptotics of Γ. For example, G is word-hyperbolic iff X is a Gromov-hyperbolic metric space, and the Gromov boundary of X turns out to be equivariantly homeomorphic to the Gromov boundary of G.

The asymptotic topological theory of splittings over finite subgroups turned out to be rather simple, eventually. Given our group G as above, one considers the space of ends of Γ, or, rather, its cardinality $e(G)$: a classical argument by Hopf quickly yields the following classification:

$e(G) = 0$**.** G is a finite group;

$e(G) = 1$**.** G is infinite, any compactification of G has connected boundary;

$e(G) = 2$**.** G is virtually-cyclic;

$e(G) = \infty$**.** Actually, $e(G) = 2^{\aleph_0}$, as the endspace of Γ can be shown to be perfect.

In the latter two cases, Stallings' theorem tells us that G splits over a finite subgroup. In the case when G is finitely presented, the accessibility theorem of Dunwoody further guarantees that G decomposes as a finite graph of groups with finite edge-groups and one-ended vertex groups.
In the context of this work it is important to stress the topological interpretation of a group being one-ended. It is not hard to show that $e(X)$ is a quasi-isometry invariant. Therefore, $e(G) = e(X)$ for all geodesic metric spaces admitting a geometric group action by G. If now $\hat{X}$ is *any* compactification of X containing X as an open subspace – call such compactifications *good*, – then the boundary $\partial_\infty X = \hat{X} \setminus X$ is necessarily connected, being the intersection of a descending chain of connected compact subset of $\hat{X}$. Thus, in order for G to have more than one end (and, consequently, to split over a finite subgroup) it is necessary and sufficient that some geodesic space admitting a geometric action by G have a good compactification with disconnected boundary.

To summarize, in order to employ topological methods for the study of relations between asymptotic/coarse properties of G and its various splittings it will be enough, in view of Dunwoody accessibility, to consider one-ended finitely presented groups G.
At this point it is noteworthy to emphasize that the various techniques employed by different authors to provide "geometric proofs" of Stallings' theorem, as well as other splitting results, have all utilized the same result by Bass and Serre: a group G splits if and only if it acts on a directed simplicial tree T without a global fixed vertex and without a global fixed end.
The proof of this theorem provides a tool allowing us, once we are able to control the edge-stabilizers of T, to have precise information regarding the nature of the splittings arising from T.

Unfortunately, the approaches that worked for the relatively-hyperbolic case ([Bow98, Gur05]) turned out to be hard to apply to the CAT(0) setting, because the techniques used depend strongly on the local connectivity of the boundaries involved, as well as on special dynamic properties that the action of a relatively-hyperbolic group on its canonical boundary has (for example, one question – still open –, asks when is the action of a group on a CAT(0) boundary minimal). The few known examples of CAT(0) groups whose boundaries seem to portray information regarding known splittings of these groups are also known to have non-locally-connected multiple distinct boundaries (see [MR99, MR01, CK00]), and the methods that were originally developed for (relatively) hyperbolic groups fail. Therefore, it seemed a more immediate goal to try and pin-down those connectivity properties of a CAT(0) group which turn up in *all* of its boundaries. In order to do so, we turned back to analyzing the end structure of G.
The main problem realizing the same approach for splittings of one-ended groups over infinite subgroups has been the absence of a good combinatorial object whose structure expresses what one would like to consider as the "relative end structure" of the pair (G, H).
While in the case when H is the trivial subgroup of G and the pair (G, H) is multi-ended one naturally expects (from the point of view of coarse geometry) the "end structure" of the pair (G, H) to be the same as that of a tree, when dealing with infinite subgroups H this is not what one gets. Instead, in his thesis ([Sag95]) Sageev has shown that what one usually gets is an action on a higher-dimesional analog of a tree – namely, a *cubing*, or *non-positively curved cube complex*. Still, the action one obtains is non-trivial in a sense similar to the non-triviality condition of Bass and Serre, so that one may hope to find an equivariant retraction of this cube complex onto a tree in order to achieve a splitting. The main problems with this approach are:

- Given the subgroup H, the "Sageev cubing" is not uniquely determined by H; even slight modifications in the choices involved in Sageev's construction may change the dimension of the complex drastically;

- Unlike the case of finite subgroups, it is not at all clear whether two different Sageev cubings of the same pair (G, H) have the same "trace" on the boundary of a given compactification of X.

- Unlike the tree constructed for the proof of Stallings theorem, or the end com-

pactification of Γ, there is no notion of boundary for cubings that will be coarse enough to enable comparison of coarse properties of Sageev cubings with coarse properties of Γ (or of any geodesic space X admitting a geometric action by G).

- There is no known bound on the dimension of the Sageev cubings.

Of these, the main obstacles seemed to be the second and third. Thus, the first part of this work deals with constructing a combinatorial boundary for cubings and with its properties. This involves deepening the analysis of Roller's duality between cubings (in their incarnation as discrete median algebras) and ω-dimensional poc-sets (surveyed in section 2) by investigating hierarchical relations among the distinct components of the double-dual of a cubing (section 3.2). In honor of Roller's pioneering work on this duality, we call the resulting structure by the name *the Roller boundary of a cubing* (defined in 3.3).

Because of the abundance of open questions regarding the Sageev cubings of a group pair, the main test-case for our construction in this thesis was chosen to be less problematic. We assume our group G acts geometrically on a CAT(0) space X, with X containing a fixed G-invariant *halfspace system*. The idea had been – before any attempts at coarse geometry are made – to apply the Roller boundary to better tailored objects, like hyperplanes in Euclidean or Hyperbolic space, or walls in the Davis-Moussong complex of a Coxeter group. We have borrowed the notion of a wall and of a halfspace from these geometric examples, obtaining the following definitions: a *halfspace* in a CAT(0) space X is a convex open subset $h \subset X$ such that $h^* = X \smallsetminus \overline{h}$ is also convex; the *wall* $W(h)$ of a halfspace h is then defined to be $W(h) = \overline{h} \cap \overline{h^*}$; a *halfspace system* is then a set $\mathcal{H}$ of halfspaces which is closed under the operation of replacing h by h^*, and satisfying some obvious discreteness conditions, like, for example, that of any point in X having a neighbourhood meeting only finitely-many walls of $\mathcal{H}$.

It is a crucial point in this work that one can use Sageev's construction to show that a halfspace system in CAT(0) space naturally defines a cubing $\mathcal{C}(\mathcal{H})$. If the halfspace system $\mathcal{H}$ is G-invariant, then many well-known examples (see [Sag95] for a general discussion, for Coxeter groups see [NR03, Wil98], for small cancellation groups see [Wis04]) demonstrate the importance of knowing when is the action of G on this cubing cocompact. In this work, we answer this question in terms of properties of a canonical map $\rho_{\mathcal{H}}$ we have defined from the CAT(0) boundary of X into the Roller boundary of $\mathcal{C}(\mathcal{H})$, thus showing that this question is of a coarse nature (see theorem 6.1).

Apart from this main result, we show several other uses for the map ρ. In particular, we show that ρ induces a stratification of $\partial_\infty X$ with respect to the cone topology (proposition 4.24), modeled on the hierarchical structure of the Roller boundary $\Re\mathcal{H} = \Re C(\mathcal{H})$, and that the same stratification is inherited by convex subspaces of X (proposition 4.28). Under additional geometric assumptions on $\mathcal{H}$, which are satisfied by CAT(0) cube complexes as well as by Davis-Moussong complexes of Coxeter groups of finite rank, we provide a formula calculating the closure of a stratum (theorem 5.5), prove the injectivity of ρ on subsets of $\partial_\infty X$ which are π-discrete in the angular metric (corollary 5.3), and show how the decomposition into strata computes the Tits path-components

of $\partial_\infty X$ (proposition 5.8 and theorem 5.13). In particular, we show how our tools work for several known examples of groups and their CAT(0) boundaries, such as the Croke-Kleiner example.

Acknowledgements. some of the results herein were obtained, though in a slightly weaker form, as parts of the author's Ph.D. thesis. The generous support of the Technion IIT is gratefully acknowledged. Many thanks go to my scientific advisors Michah Sageev and Bronislaw Wajnryb, for giving me all the support and guidance I could ever hope for.

2 Preliminaries: Sageev-Roller duality.

A cubing is a piecewise-Euclidean simply-connected cell complex $\mathcal{C}$ satisfying the following requirements: all cells of $\mathcal{C}$ are standard Euclidean cubes; all attaching maps are Euclidean isometries; no two k-faces of the same d-cube are attached to each other, for all $0 \leq k \leq d$; all links are (simplicial) flag-complexes (Gromov's so called *link condition* for non-positive curvature).
Given a cubing $\mathcal{C}$, it is possible to make $\mathcal{C}$ a length-metric space using the identification of each d-dube with the unit Euclidean d-cube. By a theorem of Bridson ([BH99], theorems I-7.19,I-7.50) and the so-called Cartan-Hadamard theorem ([BH99] theorem II-4.1, combined with II-5.20), if $\mathcal{C}$ is finite-dimensional, or if $\mathcal{C}$ is locally finite, then the induced path pseudo-metric on $\mathcal{C}$ is a complete CAT(0) metric and every cube of $\mathcal{C}$ is isometrically embedded in $\mathcal{C}$.

2.1 The halfspace structure of a cubing.

In a cubing $\mathcal{C}$, the notion of a wall/halfspace arises naturally in combinatorial form. Two edges in a cubing are said to be parallel, if there exists a 2-cube containing them as opposite edges; one extends this notion of parallelism over the 1-skeleton of $\mathcal{C}$ by taking its transitive closure; next, given a d-dimensional cube Q of $\mathcal{C}$, one may divide its 1-skeleton into d disjoint classes of edges which are parallel in Q, and define a *midplane* of Q with respect to a class E (computed in Q – not in $\mathcal{C}$) to be the convex hull in Q of the set of midpoints of edges in E. Note that if M is a midplane of a cube Q, and Q' is a face of Q, then $M \cap Q'$ is a midplane of Q' whenever this intersection is non-empty. If Q_1, Q_2 are adjacent cubes, and Q' is their maximal common face, then the midplanes M_i of Q_i are said to be compatible if and only if the sets $M_i \cap Q'$ are non-empty and equal. Once again extending the compatibility relation transitively, we say that the union of a compatibility class of midplanes in $\mathcal{C}$ is *a wall*. The following is proved in Sageev's thesis:

Theorem 2.1 ([Sag95], theorems 4.10, 4.11) *Suppose $\mathcal{C}$ is a cubing and W is a wall of $\mathcal{C}$. Then:*

1. *W is itself a cubing;*
2. *W does not self-intersect – that is: the intersection of W with any cube of $\mathcal{C}$ either equals the empty set or a unique midplane of that cube.*

3. *W separates $\mathcal{C}$ into the union of precisely two connected componets, whose common boundary in $\mathcal{C}$ equals W – these components are called* the halfspaces determined by W. ■

Note: One may also extend the discussion in the proof of theorem 4.13 of [Sag95] or use Roller's duality results in order to show that the halfspaces of a cubing are convex.

2.2 Duality of poc-sets and cubings.

It has been Sageev's major discovery that a cubing may be fully reconstructed from the nesting patterns of its halfspaces. We present the construction in the form introduced by Roller in [Rol98], as an application of a duality theory for median algebras.

Definition 2.2 (poc-set, nesting, transversality) *A poc-set $(H, \leq, *)$ is a partially-ordered set $(H, \leq)$ with a minimum element 0 and an order-reversing involution $h \mapsto h^*$ satisfying the requirement that for all $h \in H$, if $h \leq h^*$ then $h = 0$.*

- *the elements $0, 0^*$ are said to be the* trivial *elements of H, while all other elements of H are* proper.
- *the poc-set $(H, \leq, *)$ is said to be* discrete*, if, for every pair of proper elements $a, b \in H$, the interval $[a, b] = \{h \in H \mid a \leq h \leq b\}$ is finite.*
- *two elements $h, k \in H$ are said to be nested (resp. transverse), – denoted here with $h \| k$ (resp. $h \pitchfork k$) – if one (resp. none) of the relations $h \leq k$, $h^* \leq k$, $h \leq k^*$, $h^* \leq k^*$ holds. A subset $S \subseteq H$ is nested (resp. transverse) if all its elements are pairwise nested (resp. transverse).*
- *a poc-set $(H, \leq, *)$ is said to have dimension at most d if every transverse subset of H has at most d elements. H is said to be ω-dimensional, if it contains no infinite transverse subset.*

A basic observation by Sageev is that the halfspace system $\mathcal{H}$ of a cubing $\mathcal{C}$ is a discrete poc-set with respect to inclusion and the complementation operator defined by $h^* = \mathcal{C} \smallsetminus \overline{h}$.

2.2.1 Ultrafilters.

Suppose now we are given a discrete poc-set H, and we wish to construct a cubing $\mathcal{C}$ whose natural halfspace system $\mathcal{H}$ is poc-isomorphic to H (with the obvious definition of a poc-morphism as a $(*)$-equivariant morphism of posets). One constructs a dual space for H:

Definition 2.3 *Suppose $(H, \leq, *)$ is poc-set. An ultrafilter α on H is a subset of H satisfying:*

(**UF1**) *for all $h \in H$, either $h \in \alpha$ or $h^* \in \alpha$, but not both;*

(**UF2**) *for all $h, k \in \alpha$, the relation $h \leq k^*$ is prohibited.*

The space of all ultrafilters on H will be denoted by H°.
A collection $\alpha \subset H$ satisfying (**UF2**) *is called a* filter base.

Remark 2.4 *Zorn's lemma implies that any filter base is contained in an ultrafilter.*

Remark 2.5 *H° inherits a topology from 2^H (the Tychonoff topology). In fact, H° is a Stone space with respect to this topology.*

One motivation for considering ultrafilters on a poc-set is that when H is the natural halfspace system of a cubing $\mathcal{C}$, the vertices of $\mathcal{C}$ may be mapped naturally into H° by sending every vertex v to the ultrafilter π_v of all halfspaces containing v. It is clear that ultrafilters arising in this way satisfy a *descending chain condition*:

Definition 2.6 *An ultrafilter α on a poc-set is* principal*, if it satisfies the* descending chain condition *(DCC): if $(h_n)_{n=1}^\infty$ is a descending chain of elements in α in the sense that $h_{n+1} \leq h_n$ for all natural n, then $h_{n+1} = h_n$ for all but finitely-many values of n.*

Roller considered a stronger condition for the purpose of distinguishing the ultrafilters of the form π_v:

Definition 2.7 *An ultrafilter α on a poc-set H is* well-founded*, if, for every $a \in \alpha$, α contains only finitely many elements h satisfying $h \leq a$.*

2.2.2 Almost-equality and (re)constructing cubings.

Since H is discrete, it makes sense to consider the following metric on H° (infinite values are allowed):

$$\Delta(\alpha, \beta) = \frac{1}{2}\big|\alpha \triangle \beta\big| = |\alpha \smallsetminus \beta| = |\alpha \cap \beta^*| \ , \tag{1}$$

When $\Delta(\alpha, \beta)$ is finite, we say that α, β are *almost equal*. For example, if H is the halfspace system of a cubing, then the ultrafilters π_v described above all lie in the same almost-equality class.
The metric Δ allows one to construct a graph $\Gamma = \Gamma(H)$ with vertex set H° and with $\alpha, \beta \in H^\circ$ joined by an edge iff $\Delta(\alpha, \beta) = 1$. It is clear that the components of Γ correspond to almost-equality classes of H°. Henceforth, we will use Γ_Σ to denote the component of an almost-equality class Σ in the graph Γ. A more detailed study of these graphs will be done in the next section, while here we focus on how they are used for constructing cubings.

Sageev (in a special case – cite[Sa]) and Roller [Rol98] showed that each of the graphs Γ_Σ is the 1-skeleton of a cubing:

1**-skeleton.** An important element of their observation was that the restriction of Δ to a class Σ coincides with the combinatorial metric Σ induced from Γ_Σ. This facilitated the following inductive construction of a cubing C_Σ whose 1-skeleton is $C^1_\Sigma = \Gamma_\Sigma$.

d-skeleton. Given $d \geq 2$ and a piecewise-Euclidean cubical $(d-1)$-complex C_Σ^{d-1} (consisting of unit cubes), one glues a unique isometric copy of $[0,1]^d$ onto every occurrence of $\partial[0,1]^d$ in C_Σ^{d-1}. The resulting complex is C_Σ^d.

Thus, if one wants to reconstruct a cubing C from its halfspace system, one 'only' needs to select the right almost-equality class. Sageev and Roller have shown that the almost-equality class – call it Π – of the ultrafilters π_v defined above is the right choice, in the sense that the vertex map $v \mapsto \pi_v$ induces an isomorphism of cubings $C \to C_\Pi$.

In the general situation (when H is provided abstractly) it is harder to select the right almost-equality class. And what is the meaning of 'right' anyway? In [Gur06] the author shows that both the sets of principal ultrafilters and of well-founded ultrafilters form unions of almost-equality classes. It can also be shown (see [Rol98], proposition 9.4) that every well-founded almost-equality class is Tychonoff-dense in H°, and one easily sees that H can be recovered from the corresponding cubing (although the different cubing arising in this way may not be isomorphic). Thus, knowing H is the halfspace system of a cubing C, provides a way for associating a canonical dense almost-equality class of H° with H the cubing C.
In general though, one needs to know more about the poc set H in order to select a 'nice' canonical class. In [Gur06] it is shown that a canonical class of principal ultrafilters exists in either of the following cases:

H has a principal ultrafilter of finite degree in Γ. Corollary 2.9 shows that such a vertex of Γ is a well-founded ultrafilter, and a computation shows that *in this case* there can be only one well-founded class.

H is ω-dimensional. In this case there is only one almost-equality class of principal ultrafilters, which is then also well-founded (corollary 3.3).

In either case, we denote the obtained canonical class by Π, and call it the *principal class* of H°. In the geometric application discussed in this paper the first condition will be fulfilled automatically, while for some of the applications we will have to invoke the additional assumption of ω-dimensionality on H.

2.2.3 Generalities about the structure of the graph $\Gamma(H)$.

Our analysis of cubings will require some technical details. The following observations relating to the local structure of the graph Γ will be used frequently, and appear already in [Sag95].

Given an edge of Γ with endpoints $\alpha, \beta \in H^\circ$, it is clear that β has the form

$$[\alpha]_a = (\alpha \smallsetminus \{a\}) \cup \{a^*\} \tag{2}$$

for some $a \in \alpha$. It is readily seen that β being an ultrafilter implies $a \in \min(\alpha)$: by this we mean that a has to be a minimal element in α, where α is considered with the ordering induced from H. Thus

Lemma 2.8 *The set of neighbours of any given $\alpha \in H^\circ$ is parametrized by* $\min(\alpha)$.

For example, it immediately follows that –

Corollary 2.9 *If $\alpha \in H^\circ$ is principal and has finite degree then α is well-founded.*

Another case of interest arises when α is the vertex of a d-dimensional cube Q: in this case, if $a_1, \ldots, a_d$ are the minimal elements of α corresponding to its neighbours in Q, then, for every pair of distinct i, j one must have

$$[[\alpha]_{a_i}]_{a_j} = \left[[\alpha]_{a_j}\right]_{a_i} \tag{3}$$

which implies $a_i \pitchfork a_j$ for all such pairs. We obtain –

Lemma 2.10 *For any $\alpha \in H^\circ$, there is a one-to-one correspondence between transverse subsets of $\min(\alpha)$ of size d and the d-dimensional cubes Q arising from the Sageev construction and satisfying $\alpha \in Q^0$.*

In particular, if H is a d-dimensional discrete poc-set, then C_Σ is at most d-dimensional, for every almost-equality class Σ of H°; if H is discrete and ω-dimensional, then C_Σ contains no infinite dimensional cube. The converse is, unfortunately, false.

2.2.4 Convexity structure and topology on H°.

Here we provide a brief account on convexity results from [Rol98] that are used in this paper.

Recall that 2^H – the power set of H – carries a natural *median operation*:

$$\mathbf{med}(\alpha, \beta, \gamma) = (\alpha \cap \beta) \cup (\beta \cap \gamma) \cup (\alpha \cap \gamma), \tag{4}$$

making 2^H into a *median algebra*. It is straightforward to verify that H° is closed under this operation. This allows one to define a notion of convexity in H°: given $\alpha, \beta \in H^\circ$, *the interval with endpoints* α, β is defined to be

$$[\alpha, \beta] = \left\{\gamma \in H^\circ \,\middle|\, \gamma = \mathbf{med}(\alpha, \beta, \gamma)\right\}; \tag{5}$$

a subset Σ of H° is said to be *convex*, if it contains $[\alpha, \beta]$ for all $\alpha, \beta \in \Sigma$. A convex subset of H° is a *halfspace*, if its complement is also convex.

When H is a discrete poc-set, the following subsets of H° are *the halfspaces* of H° –

$$S_h = \left\{\alpha \in H^\circ \,\middle|\, h \in \alpha\right\}, \tag{6}$$

where h ranges over the whole of H. It is shown in [Rol98] that a closed subset of H° is convex iff it is the intersection of a family of halfspaces.

In any median algebra – in particular, in H° – we have the following result:

Theorem 2.11 (Helly's Theorem ([Rol98])) *Let $C_1, \ldots, C_m$ be pairwise intersecting convex subsets of a median algebra M. Then $\bigcap_{i=1}^m C_i$ is non-empty.*

The common application of Helly's theorem in this work will be the following:

Corollary 2.12 *Suppose $B \subset H$ is a finite filter-base in H, and Σ is an almost-equality class of H° such that $\Sigma \cap S_b$ is non-empty for all $b \in B$. Then $\Sigma \cap \bigcap_{b\in B} S_b$ is non-empty.*

The proof of the corollary relies on the following simple fact:

Lemma 2.13 *Every almost-equality class of H° is a convex subset of H°.*

Proof : Fix an almost-equality class Σ. In order to verify our claim, we need to take any $\sigma \in \mathcal{H}^\circ$, a pair of ultrafilters $\alpha, \beta \in \Sigma$, and check that $\mu = \mathbf{med}(\alpha, \beta, \sigma)$ lies in Σ. Let us write

$$\mu = (\alpha \cap \beta) \cup (\alpha \cap \sigma) \cup (\beta \cap \sigma), \tag{7}$$

and observe that

$$\mu \smallsetminus \alpha \subseteq (\beta \cap \sigma) \smallsetminus \alpha \subseteq \beta \smallsetminus \alpha. \tag{8}$$

Since the right-hand side is finite, we are done. ■

Proof of corollary 2.12: The proof is straightforward now: the family $\{\Sigma\} \cup \{S_b\}_{b\in B}$ satisfies the assumptions of Helly's theorem, so its total intersection is non-empty. ■

We have already referred to closed subsets of H° and the Tychonoff topology. Let us make this precise. We topologize H° as a subspace of 2^H with the product topology. We will refer to this topology as the Tychonoff topology on H°. It is easy to see that the family of all subsets of H° of the form

$$S_h = \big\{\alpha \subseteq H \,\big|\, h \in \alpha\big\} \tag{9}$$

is a sub-base for the Tychonoff topology H°. Thus, if $A \subseteq H$ is any finite set, then $V(A) = \bigcap_{a\in A} S_a$ is a (possibly empty) basic open set (for H°); by Helly's theorem, for $V(A)$ to be non-empty it is necessary and sufficient that no relation of the form $a \leq b^*$ holds for $a, b \in A$, so the family of all $V(A)$ with A ranging over all finite filter-bases of H constitutes a basis for the Tychonoff topology on H°.

The space H° is obviously Hausdorff. Moreover, it is totally disconnected, since the sets S_h are all clopen. To see that H° is compact it is enough to verify it is closed in 2^P (we leave this as an exercise to the reader). It follows that

Proposition 2.14 *H° is a Stone space. Moreover, if $\mathcal{F}$ is a family of pairwise-intersecting closed convex subsets of H°, then $\bigcap \mathcal{F} \neq \varnothing$.*

2.2.5 Functorial properties.

Suppose now that H_1, H_2 are poc-sets. A morphism of poc-sets is, by definition, a $(*)$-equivariant morphism of posets $f : H_1 \to H_2$ satisfying $f(0) = 0$. It is easy to see that, given a poc-set H, the space H° can be naturally identified (using characteristic functions) with the space of poc-morphisms into the 2-element poc-set. As a result, given f as above, there is a natural pullback map $f^\circ : H_2^\circ \to H_1^\circ$. For our purposes it will be important to state the following facts from [Rol98], sections 3 and 4:

Proposition 2.15 *The assignment $H \mapsto H^\circ$, $f \mapsto f^\circ$ is a contravariant functor of the category of poc-sets with poc-morphisms into the category of Stone median algebras with continuous median-preserving maps.*

The pullback in the poc-set category is exact in the following sense:

Proposition 2.16 *Let $f : H_1 \to H_2$ be a morphism of poc-sets.*

- *f° is injective iff f is surjective;*
- *f° is surjective iff f is an embedding.*

3 A boundary for ω-dimensional cubings

Throughout this section, $(H, \leq, *)$ is a discrete ω-dimensional poc-set. Our initial goal will be to establish some elementary facts about how the combinatorics of infinite descending chains is related to the structure of a given almost-equality class and its Tychonoff closure. We first study the principal class.

3.1 Characterizing the principal class.

Consider the following way of constructing an ultrafilter:

Proposition 3.1 *Let A be a maximal transverse subset of H. Then the collection*

$$\pi_A = \left\{h \in H \,\middle|\, \exists_{a\in A}\ a \leq h\right\} \cup \left\{h \in H \,\middle|\, \exists_{a\in A}\ a^* < h\right\} \tag{10}$$

is a principal ultrafilter on H. Moreover, every principal ultrafilter is of this form.

Proof : First let us show $\pi_A \in H^\circ$. We need to show three things:

- For any $h \in H$, either $h \in \pi_A$ or $h^* \in \pi_A$.

  Given $h \in H$, if neither h nor h^* are greater than an element of $A \cup A^*$, it means that $A \cup \{h\}$ is a transverse set, contradicting the maximality of A as such a set.

- For any $h \in H$, if $h \in \pi_A$ then $h^* \notin \pi_A$.

  If *both* h and h^* lie in π_A, then there are $a, b \in A \cup A^*$ such that $a \leq h$ and $b \leq h^*$. This implies $a \leq h \leq b^*$, and the transversality of A then gives $a = b^*$, resulting in $a = h$. Thus, both $h = a$ and $h^* = a^*$ lie in π_A, which is impossible, by the construction of π_A.

- If $h, k \in \pi_A$, then $h \not\leq k^*$.

  Suppose the contrary, and let $a, b \in A \cup A^*$ be such that $a \leq h$ and $b \leq k$. Then we have

  $$a \leq h \leq k^* \leq b^*, \tag{11}$$

  meaning again that $b^* = a$, and hence $h = k^*$, finally implying $k, k^* \in \pi_A$, which we have just shown to be impossible.

In order to prove π_A is principal, assume $\eta = (h_n)_{n=1}^{\infty}$ is a descending chain in α, and we have to show it is eventually constant. Indeed, for each n find $a_n \in A \cup A^*$ such that $a_n \leq h_n$. Since $A \cup A^*$ is finite, the a_n may be taken to be constant – denote them by a. Thus $h_n \in [a, h_1]$ for all n. Since H is discrete, the h_n stabilize.

Conversely, suppose α is a principal ultrafilter, and consider the set $\min(\alpha)$ of all minimal elements of α. Since α is principal, it is clear that every $h \in \alpha$ has some $a \in \min(\alpha)$ satisfying $a \leq h$. In particular, $\min(\alpha)$ is non-empty.
We will now show $\min(\alpha)$ contains a maximal transverse set A of H; in particular, α contains π_A, which is an ultrafilter. Since no two distinct ultrafilters contain each other, we will conclude $\pi_A = \alpha$, completing the proof.
We now construct the required set A by induction. Start from any element $a_1 \in \min(\alpha)$, and let $A_1 = \{a_1\}$. If A_1 is a maximal transverse subset of H, then we are done – else, proceed by induction as follows: suppose $A_i \subset \min(\alpha)$ is a given transverse subset which is not maximal in H, then there is an element h of H with $h \pitchfork A_i$, and then there is an element $a \in \min(\alpha)$ with either $a \leq h$ or $a \leq h^*$; in any case, $a \pitchfork A_i$, and we set $A_{i+1} = A_i \cup \{a\}$. Since H contains no infinite strictly-ascending chain of transverse subsets, this process must stop, producing a subset A of $\min(\alpha)$ which is a maximal transverse subset of H, as desired. ■

Corollary 3.2 *$\alpha \in H^\circ$ is principal iff* $\min(\alpha)$ *contains a maximal transverse subset A of H, in which case $\alpha = \pi_A$.*

An important consequence of this is:

Corollary 3.3 *The principal ultrafilters on H form an almost-equality class in H°.*

Proof : Suppose $\alpha = \pi_A$ and $\beta = \pi_B$ are principal ultrafilters, and consider $h \in \alpha \smallsetminus \beta$: we have $a \leq h$ for some $a \in A \cup A^*$, but $h \notin \beta$; since B is a maximal transverse set, $B \cup \{h\}$ is not transverse, there is a $b \in B$ such that the pair $\{b, h\}$ is nested. The assumption $h \notin \pi_B$ then forces either $h < b$ or $h \leq b^*$, so that h belongs to $[a, b] \cup [a, b^*]$. Thus, we have shown that $\beta \smallsetminus \alpha$ lies in the union of all intervals of the form $[a, b]$ with $a \in A \cup A^*$ and $b \in B \cup B^*$ whenever those are defined; since H is discrete and both A and B are finite, we are done. ■

Henceforth, we shall denote the set of principal ultrafilters on H by $\Pi(H)$, or simply by Π, when H is understood from the context.

Corollary 3.4 *Π is dense in H°. Moreover, it is the only almost-equality class of H° with this property.*

Proof : Let us prove Π is dense. Let $V(A)$ be an arbitrary basic open set in H°, with A a finite filter base in H. Then the corollary (2.12) to Helly's theorem simply states that $V(A) \cap \Pi$ is non-empty, provided we can show that $\Pi \cap S_a$ is non-empty for all $a \in A$. In fact, $\Pi \cap S_a \neq \varnothing$ for any proper $a \in H$: given a, let A be any maximal transverse subset containing a. Then π_A lies in $\Pi \cap S_a$, by proposition 3.1.
We defer the proof of the second part until after we have studied some combinatorial properties of almost-equality classes. ■

3.2 Almost-equality of ultrafilters

Given an almost-equality class Σ, recall that the metric Δ on H° restricts on Σ to the combinatorial metric of the graph Γ_Σ (this result, initially due to Sageev in a special case, was proved in this generality, independently in [Rol98], [Nic04] and [CNar] with slight variations of context).
Thus, a vertex path $(\sigma_0, \ldots, \sigma_d)$ from $\alpha \in \Sigma$ to $\beta \in \Sigma$ must satisfy $d \geq \Delta(\alpha, \beta)$, and may be chosen so that equality holds. Let us recall that for each $i \in \{1, \ldots, d\}$ there are $h_i \in \min(\sigma_i)$ such that $\sigma_{i+1} = [\sigma_i]_{h_i}$. It will be convenient to consider such a vertex path as a chain (of length d) of transformations of α into β, performed in stages. For this reason, we shall also refer to the (directed) edges of Γ_Σ as *elementary moves*.

3.2.1 Parallel/Transverse decomposition.

Let us fix an almost-equality class Σ of H°. Obviously, some elements of H have nothing to do with separation in Σ:

Definition 3.5 *Suppose $h \in H$ and Σ is an almost-equality class in H°. We say h* is transverse to Σ *(denoted by $h \pitchfork \Sigma$), if both $S_h \cap \Sigma$ and $S_{h^*} \cap \Sigma$ are non-empty. Otherwise we shall say that h is* parallel to Σ. *We shall keep the following notation:*

$$T(\Sigma) \quad = \quad \{h \in H \,|\, h \pitchfork \Sigma\} \cup \{0, 0^*\}. \tag{12}$$

With the induced ordering, $T(\Sigma)$ is a discrete ω-dimensional poc set in its own right. Its less abstract counterpart is the poc-set

$$H_\Sigma = \left\{ S_h \cap \Sigma \,\middle|\, h \in H \right\} \tag{13}$$

viewed as a sub poc-set of the power set of Σ. Note that H_Σ is the image of H under the *restriction map*

$$\left\{ \begin{array}{rcl} r_\Sigma : H & \to & H_\Sigma \\ h & \mapsto & S_h \cap \Sigma\,, \end{array} \right. \tag{14}$$

mapping $T(\Sigma)$ onto H_Σ, and sending elements parallel to Σ to the trivial elements of H_Σ. Obviously, r_Σ is a morphism of poc-sets.

Lemma 3.6 *Assume $a \pitchfork \Sigma$. Then there exists an ultrafilter $\gamma \in S_a \cap \Sigma$ such that $a \in \min(\gamma)$. Conversely, if $\gamma \in \Sigma$ and $a \in \min(\gamma)$, then $a \pitchfork \Sigma$.*

Proof : Let $\gamma \in S_a \cap \Sigma$ and $\gamma^* \in S_{a^*} \cap \Sigma$. Then there is a sequence of Σ-edges from γ to γ^* – in particular, a is contained in the minimal set of some element of this sequence. For the converse, simply consider $[\gamma]_a$: this is an element of Σ contained in S_{a^*}. Thus, neither $S_a \cap \Sigma$ nor $S_{a^*} \cap \Sigma$ is empty. ■

Corollary 3.7 *Suppose $\sigma \in \Sigma$ contains an infinite descending chain $c = (c_n)_{n=1}^\infty$. Then $\Sigma \subseteq S_{c_n}$ for all n.*

Proof : Fix n, let $c = c_n$, and suppose $c \pitchfork \Sigma$. Lemma 3.6 then implies there exists $\xi \in \Sigma$ containing c in its minimal set. In particular, $\sigma \smallsetminus \xi$ contains the infinite set

$\{c_m \,|\, m > n\}$, contradicting the fact that σ and ξ are almost-equal. Thus, c_n is parallel to Σ, and being an element of $\sigma \in \Sigma$, necessitates $\Sigma \subseteq S_c$, as desired. ■

From this we obtain the topological characterization of the principal class:
Proof of cor. 3.4: We have shown already that the principal class $\Pi(H)$ is dense in H°. Now, suppose that Σ is non-principal. Then some $\sigma \in \Sigma$ contains an infinite descending chain, implying there exist proper elements $h \in H$ satisfying $\Sigma \subseteq S_h$. In particular, Σ does not intersect the non-empty open subset S_{h^*} of H°. ■

Lemma 3.8 *Suppose $a, b \pitchfork \Sigma$. Then, $S_a \cap \Sigma \subseteq S_b \cap \Sigma$ holds if and only if $a \leq b$. Thus, r_Σ restricts to an isomorphism between the poc-sets $T(\Sigma)$ and H_Σ. In particular, H_Σ is a discrete ω-dimensional poc-set.*

Proof : One direction is trivial. For the other, we assume both $S_a \cap \Sigma \subseteq S_b \cap \Sigma$ and $a \not\leq b$ hold, and seek to derive a contradiction.
Among all pairs $(\alpha, \beta) \in \Sigma \times \Sigma$ satisfying the condition

$$(*) \quad a \in \min(\alpha) \ \wedge \ b^* \in \min(\beta) \tag{15}$$

let us fix one for which $n = \Delta(\alpha, \beta)$ is minimal. Note that lemma 3.6 guarantees the set of all pairs satisfying $(*)$ is non-empty.

Let $(\sigma_0 = \alpha, \ldots, \sigma_n = \beta)$ be a chain (of elementary moves) of length n. We define

$$i \ = \ \max\left\{t \in \{0, \ldots, n\} \,\middle|\, \forall_{s \leq t}\ a \in \sigma_s\right\} \tag{16}$$

$$j \ = \ \min\left\{t \in \{0, \ldots, n\} \,\middle|\, \forall_{s \geq t}\ b^* \in \sigma_t\right\}. \tag{17}$$

Since $a \in \sigma_i$, our initial assumption implies $b \in \sigma_i$, so that $i < n$. Note now that, by the definition of i, $a^* \in \sigma_{i+1}$, so that $a \in \min(\sigma_i)$. Thus, if i were greater than zero, then the chain $(\sigma_i, \ldots, \sigma_n)$ satisfies $(*)$ – a contradiction to n being the minimal length of such a chain. Similarly, one proves $j = n$, and we have

$$t = 1, \ldots, n \quad \Rightarrow \quad a^* \in \sigma_t \tag{18}$$

$$t = 0, \ldots, n-1 \quad \Rightarrow \quad b \in \sigma_t. \tag{19}$$

Note that $a^* \in \min(\sigma_1)$ and $b \in \min(\sigma_{n-1})$.
Assume $a^* \in \min \sigma(t)$, and let us show $a^* \in \min(\sigma_{t+1})$. By definition of an elementary move, there is an element $c \neq a^*$, $c \in \min(\sigma_t)$ such that $\sigma_{t+1} = [\sigma_t]_c$. If a^* is not minimal in σ_{t+1}, then it means $c^* < a^*$, which implies $c < a$. But then $c \in \sigma_t$ implies $a \in \sigma_t$, which is impossible, as $a^* \in \sigma_t$.
Thus, we have shown $a^* \in \min(\sigma_t)$ for all $t = 1, \ldots, n$. In particular, the ultrafilter $\sigma = \sigma_{n-1}$ contains both a^* and b as minimal elements, and $a^* \not\leq b^*$.
Recall now that $a, b \in \alpha$, so that $a \not\leq b^*$. Further, since $a^*, b^* \in \beta$, we also have $a^* \not\leq b$. Adding to this the assumption that $a \not\leq b$, we conclude $a \pitchfork b$ holds. Thus, the ultrafilter $\mu = [\sigma]_{a^*, b}$ is an element of Σ containing both a and b^*. Now, applying our assumption that $\Sigma \cap S_a \subseteq \Sigma \cap S_b$ we obtain

$$a \in \mu \ \Leftrightarrow \ \mu \in S_a \ \Rightarrow \ \mu \in S_b, \tag{20}$$

contradicting $\mu \in S_{b^*}$. ■

3.2.2 Hierarchy in H°.

restriction to a class, corestriction, projection. Our aim now is to use the fact that r_Σ restricts to an isomorphism of $T(\Sigma)$ with H_Σ for computing the Tychonoff closure of Σ in H°: intuitively it is expected that Σ arises as the principal class for each of these poc-sets, and we want to conclude that the structure of the closure of Σ in H° depends only on the structure of Σ.

Let us consider the sequence of poc-set morphisms

$$T(\Sigma) \xrightarrow{inc.} H \xrightarrow{r_\Sigma} H_\Sigma \,. \tag{21}$$

We obtain three continuous maps: a map $\iota : H^\circ \to T(\Sigma)^\circ$ dual to the inclusion, a "corestriction" map $cor^H_\Sigma : H^\circ_\Sigma \to H^\circ$ dual to r_Σ, and – by lemma 3.8 – the homeomorphism t_Σ induced (as a push forward) by the composition of r_Σ and the inclusion map. The map t_Σ enables the construction of a "restriction map"

$$res^H_\Sigma = t_\Sigma \circ \iota : H^\circ \to H^\circ_\Sigma, \tag{22}$$

Explicitly:

$$\begin{array}{rcll} res^H_\Sigma(\sigma) & = & \left\{S_h \cap \Sigma \,\middle|\, h \in \sigma\right\} & (\sigma \in H^\circ)\,, \\ cor^H_\Sigma(\bar\sigma) & = & \left\{h \in H \,\middle|\, S_h \cap \Sigma \in \bar\sigma\right\} & (\bar\sigma \in H^\circ_\Sigma)\,. \end{array} \tag{23}$$

Proposition 3.9 *Suppose Σ is an almost-equality class in H°. Then the following hold:*

1. *$res^H_\Sigma \circ cor^H_\Sigma = id_{H_\Sigma}$. In particular, cor^H_Σ is injective.*
2. *The map $pr_\Sigma : H^\circ \to H^\circ$ defined by $pr_\Sigma = cor^H_\Sigma \circ res^H_\Sigma$ is a retraction, pointwise fixing $\overline{\Sigma}$.*
3. *cor^H_Σ is a homeomorphism of H°_Σ onto $\overline{\Sigma}$.*
4. *cor^H_Σ sends almost-equality classes to almost-equality classes. In particular, the principal class of H_Σ is mapped bijectively onto Σ.*

Proof : In the process of this proof, let us suppress the H and Σ indices in the notation for the maps in concern, referring to them as *res*, *cor* and *pr*, respectively.
Claim (1.) is immediate from the explicit forms given above, and implies pr_Σ is, indeed, a retraction.

By continuity, to prove (2.) it is enough to verify that *pr* fixes Σ pointwise. For $\sigma \in \Sigma$ we compute:

$$a \in \underbrace{(cor \circ res)(\sigma)}_{\triangleq \sigma'} \quad \Leftrightarrow \quad S_a \cap \Sigma \in r_\Sigma(\sigma). \tag{24}$$

Now, there are two possibilities: if $a \pitchfork \Sigma$, then lemma 3.8 tells us that $a \in \Sigma$; if a is parallel to Σ, then its lying in the image of r_Σ means $\Sigma \subseteq S_a$, again implying $a \in \Sigma$. Thus, we have found out that σ' is contained in σ; being an element of H°, it is equal

to σ, as desired.

For (3.), note that *cor* is a bijective continuous map onto its image, with H° Hausdorff and H°_Σ compact. Thus, it is a homeomorphism onto its image. Therefore, we now have to show $Im(pr) = Im(cor)$ is actually equal to the closure of Σ, implying *cor* is a homeomorphism of H°_Σ onto $\overline{\Sigma}$. Since $\Pi(H_\Sigma)$ is dense in H°_Σ (see cor.3.4) and *cor* is continuous, it is enough to show *cor* maps the principal class $\Pi(H_\Sigma)$ into Σ. Since $pr = cor \circ res$, this is equivalent to proving *res* maps Σ *onto* $\Pi(H_\Sigma)$. In order to do this, it will be enough to show that Σ is mapped into $\Pi(H_\Sigma)$, and then to show that every chain of elementary moves beginning in $res(\Sigma)$ may be lifted through *res* to a chain of elementary moves in Σ; of course, it is enough to consider chains of length 1. Now, for any $\sigma \in \Sigma$, note that corollary 3.7 implies $res(\sigma)$ contains no infinite descending chains; thus, *res* maps Σ into the principal class of H°_Σ, as we expected. To finish our argument, let now $\sigma \in \Sigma$ be any ultrafilter, and consider $a \in H$ such that $\bar{a} = r_\Sigma(a)$ is a minimal element of $res(\sigma)$; in particular, $a \pitchfork \Sigma$. Applying lemma 3.8 gives us that a is a minimal element of σ, meaning that both the ultrafilters $[\sigma]_a$ and $[res(\sigma)]_{\bar{a}}$ exist; it is then obvious from the formula we have for *res* that

$$(*) \qquad res\left([\sigma]_a\right) = [res(\sigma)]_{\bar{a}}. \tag{25}$$

We note that this equality is true *whenever the ultrafilters involved are well-defined*, and an analogous equality holds for *cor*; this is important for the proof of (4).
Finally, we prove (4). Assume Ξ is an almost-equality class in H°_Σ. First of all, (3.) states *cor* maps Ξ into the closure of Σ in H°, so that *pr* is pointwise fixed on $cor(\Xi)$. For every $h \in H$, if $\Sigma \subseteq S_h$, then $\overline{\Sigma} \subseteq S_h$, because S_h is closed. For all $h \in H$ we have

$$(**) \quad \Sigma \subseteq S_h \Rightarrow \Xi \subseteq S_h. \tag{26}$$

Applying the same method as in the proof of (3.), in order to prove that $cor(\Xi)$ lies inside an almost-equality class in H°, we assume $\xi \in \Xi$ and $a \pitchfork \Sigma$ such that $\bar{a} = S_a \cap \Sigma \in \min(\xi)$ and we need to show that a is minimal in $cor(\xi)$: indeed, if $b \in \xi$ satisfies $b \leq a$, then the minimality of $\bar{a}$ implies $S_b \cap \Sigma = S_a \cap \Sigma$, implying $b \pitchfork \Sigma$; but then lemma 3.8 shows that $a = b$, and we have shown $a \in \min(\xi)$.
For the reverse inclusion (showing *cor* maps Ξ *onto* an almost-equality class of H°), we need to show that if a is a minimal element in $cor(\xi)$, then $S_a \cap \Sigma$ is a minimal element of ξ (which then enables us to apply $(*)$). Now, a minimal element in $cor(\xi)$ is transverse to Ξ, which, through $(**)$, implies it is also transverse to Σ. Therefore, if $S_b \cap \Sigma \subseteq S_a \cap \Sigma$ and $S_b \cap \Sigma \in \xi$, then $b \in cor(\xi)$ either satisfies $b \pitchfork \Sigma$ and then $b = a$ (as desired), or b is parallel to Σ, implying $S_b \cap \Sigma = \varnothing$; but then:

$$\begin{aligned} S_b \cap \Sigma = \varnothing \quad &\Leftrightarrow \quad \Sigma \subseteq S_{b^*} && (27)\\ &\Rightarrow \quad \Xi \subseteq S_{b^*} && (28)\\ &\Leftrightarrow \quad S_b \cap \Xi = \varnothing && (29) \end{aligned}$$

– a contradiction to $\xi \in S_b \cap \Xi$. Thus, a is a minimal element of $cor(\xi)$, and we are done. ■

For us, an important corollary of the last proposition is the following

Corollary 3.10 *Suppose an almost-equality class Ξ intersects the closure of an almost-equality class Σ. Then Ξ is contained in the closure of Σ.*

Proof : Consider once more the map $cor^H_\Sigma : H^\circ_\Sigma \to H^\circ$, and let $\alpha \in \Xi \cap \overline{\Sigma}$. By the above proposition, α lies in the image of the corestriction map. Moreover, part (4.) of the proposition implies cor^H_Σ maps the whole almost-equality class of α into $\overline{\Sigma}$. ■

Almost-equality and parallelism. We have so far considered almost-equality classes disregarding the elements of H parallel to the respective classes. This gave us a kind of "intrinsic" approach to analyzing these classes, but could not serve to provide us with ideas regarding the way these classes interact. In the last corollary we have seen that closures of almost-equality classes intersect in almost-equality classes, this resulting in a kind of hierarchy on the set of all classes. In order to study this hierarchy, we need to have a better understanding of the way how, for a given almost-equality class Σ, the poc-set $T(\Sigma)$ is embedded in H. For this purpose, the following lemma is essential:

Lemma 3.11 *Let Σ be an almost-equality class of H°, and suppose a proper element $a \in H$ satisfies $\Sigma \subseteq S_a$. Then there exists $b < a$ with the same property.*

Proof : Since a is parallel to Σ, it is not a minimal element of any $\sigma \in \Sigma$. Let us denote by P the set of immediate predecessors of a in H, and suppose there is no (proper) $b < a$ such that $\Sigma \subseteq S_b$; then every $p \in P$ is transverse to Σ. By lemma 3.6, for each such p there exists $\sigma_p \in \Sigma$ such that $p \in \min(\sigma_p)$.
Let $p, q \in P$ be distinct, and suppose they are nested. Then, since it cannot be that $p < q$ or $q < p$, we must have either $p^* < q$ or $p < q^*$. Now, if $p^* < q$ then the calculation

$$p < a \quad \Rightarrow \quad a^* < p^* < q < a, \tag{30}$$

yields a contradiction, leaving $p < q^*$ as the only possibility. Thus, there exists *no* ultrafilter $\xi \in H^\circ$ containing both p and q, which is equivalent to saying that $S_p \cap S_q$ is empty. Now consider σ_p and σ_q: we have

$$\{p, q^*\} \subset \sigma_p\,, \quad \{q, p^*\} \subset \sigma_q, \tag{31}$$

If $\sigma_0 = \sigma_p, \ldots, \sigma_k = \sigma_q$ is a chain of elementary moves – that is, $\sigma_{t+1} = [\sigma_t]_{a_t}$ for every $t \in \{1, \ldots, k\}$ – then, since none of the σ_t may contain the sets $\{p, q\}$ and $\{p^*, q^*\}$, we arrive at a contradiction, by the definition of an elementary move. Thus we have proved that P is a transverse set.
By the assumption of ω-dimensionality, P is finite. Also, P is contained in $T(\Sigma)$. Thus, there exists a principal ultrafilter $\bar{\pi} \in \Pi(H_\Sigma)$ containing $r_\Sigma(P)$ in its minimal set, and then $\pi = cor^H_\Sigma(\bar{\pi})$ is an ultrafilter in Σ having $P \subseteq \min(\pi)$. Finally, reversing all the elements of P in π one-by-one, we obtain an ultrafilter $\pi' \in \Sigma$ with the property that $a \in \min(\pi')$ – contradiction. ■

Truncation. We define a generalization of elementary moves, which we call *truncation*. The idea behind this terminology is that a non-principal ultrafilter σ may be changed into another ultrafilter containing "a smaller number" of infinite descending chains than the original by "cutting-off" all elements of σ below a certain prescribed level.

The truncation operation on $\xi \in \Sigma$ is meant to generalize the operation of reversing the orientation of a minimal element. However, while inverting a minimal element does not change the a.e.-class of an ultrafilter, inverting a whole (descending) chain of halfspaces, say, in an ultrafilter should "move" its a.e.-class down with respect to the ordering we had introduced on the set of almost-equality classes.

Lemma 3.12 (truncation) *Let $\xi \in H^\circ$ and $b \in \xi$. Then*

$$[\xi]_b = \xi \smallsetminus \{h \in H \,\big|\, h \leq b\} \cup \{h \in H \,\big|\, b^* \leq h\} \tag{32}$$

is also an element of H°. Moreover, $[\xi]_b$ is almost-equal to ξ if and only if b does not belong to an infinite descending chain of ξ.

Remark 3.13 *Suppose $\xi \in H^\circ$, and let $a, b \in \xi$ satisfy $b < a$. Then, $[[\xi]_b]_a = [\xi]_a$.*

Proof of the lemma: Let Σ be the almost-equality class of ξ. We prove the latter assertion first. If b does lie on a descending chain $\{c_n\}_{n=1}^{\infty} \subset \xi$, then, obviously, $[\xi]_b$ is not in the same almost-equality class as ξ, as $c_n \notin [\xi]_b$ for almost all n.
For the converse, it is enough to prove $b \pitchfork \Sigma$. Indeed, if b is transverse to Σ, then there exists $\sigma \in \Sigma$ with $b^* \in \min(\sigma)$. But then σ is, by the very definition of $[\xi]_b$, almost equal to $[\xi]_b$, so that $[\xi]_b \in \Sigma$, as desired.
Thus, we assume there is no infinite descending chain in ξ through b, and we need to show that $b \pitchfork \Sigma$ holds. Now, since $b \in \xi$, we have $\Sigma \cap S_b$ is non-empty (because ξ is in there), and we only need to show that $\Sigma \cap S_{b^*}$ is non-empty, too. Assume, on the contrary, the latter intersection is empty; then we have $\Sigma \subset S_b$, and we may use lemma 3.11 to inductively construct an infinite descending chain $\{c_n\}_{n=1}^{\infty}$ such that $c_1 = b$ and with the property that $\Sigma \subset S_{c_n}$ for all n. In particular, ξ – being a member of Σ – will contain this chain.

Finally, let us show $[\xi]_b$ is an ultrafilter. As before, we break the proof into three parts:

- Suppose $h \in \xi_1$, and show $h^* \notin \xi_1$:

  Assume both $h \in \xi_1$ and $h^* \in \xi_1$. For h there are two possibilities with respect to ξ:

  - If $h \in \xi$, then $h^* \notin \xi$, but $h^* \in \xi_1$ then implies $b^* \leq h^*$, giving $h \leq b$, which is impossible for $h \in \xi \cap \xi_1$.
  - If $h^* \in \xi$, then $h \notin \xi$ together with $h \in \xi_1$ imply $b^* \leq h$, and then also $h^* \leq b$, which is, again, impossible because of $h^* \in \xi \cap \xi_1$.

- Suppose $h^* \notin \xi_1$ and show $h \in \xi_1$:

  Again, consider the two possibilities for h:

  - If $h^* \in \xi$ then $h^* \leq b$, and consequently $b^* \leq h$ implying $h \in \xi_1$;
  - If $h \in \xi$ but $h \notin \xi_1$, then $h \leq b$, which means $b^* \leq h^*$, implying $h^* \in \xi_1$ – a contradiction.

- Assume $h, k \in \xi_1$ with $h < k^*$, and derive a contradiction:

  In the same manner as before, we conduct a case-study according to h belonging or not belonging to ξ.

  - If $h \in \xi$ then $k^* \in \xi$; since $k \in \xi_1$, we conclude $k^* \leq b$. In particular, $h \leq b$, so $h \notin \xi_1$ – contradiction.
  - Thus, $h^* \in \xi$ and $h \in \xi_1$, and we have $h^* \leq b$, and then also $k \leq b$. But $k \in \xi_1$ means $k \leq b$ is impossible, and we are done. ■

3.3 The Roller boundary

It seems that Roller ([Rol98]) has been the first to consider topological properties of H° in the context of a "Stone-type" duality he had discovered between median algebras and poc-sets. Therefore, in view of our cor.3.10, we define

Definition 3.14 (Roller boundary of a poc-set) *Let $(H, \leq, *)$ be an ω-dimensional discrete poc-set.*
The Roller boundary *$\Re H$ of H is the set of almost-equality classes of ultrafilters, partially-ordered by the relation*

$$\Sigma_1 \geq \Sigma_2 \quad \overset{def}{\Leftrightarrow} \quad \Sigma_1 \cap \overline{\Sigma_2} \neq \varnothing \tag{33}$$

$$\Leftrightarrow \quad \Sigma_1 \subseteq \overline{\Sigma_2}. \tag{34}$$

3.3.1 Codimension

For our purposes it will be beneficial to develop a notion of (co)dimension in $\Re H$.
Let us fix $\xi \in H^\circ$ and denote its class with Σ. Assume Σ is not the principal class.
Consider the set of all infinite descending chains in H. If $c = \{c_n\}_{n=1}^\infty$ and $d = \{d_m\}_{m=1}^\infty$, say that c *dominates* d (denoted $c \succ d$), if for every m there exists n such that $c_n \leq d_m$. We say the chains c and d are *equivalent* (denoted by $c \sim d$), if both $c \succ d$ and $d \succ c$ hold. It is immediate that the relation $(\succ)$ induces a partial ordering on the set of equivalence classes of chains, and that every chain is equivalent to any of its infinite sub-chains.
We apply the above relations to the set D_ξ of infinite descending chains which are contained in ξ. Our main interest lies with the set $D_\xi/\sim$.
Note that for any $\sigma \in \Sigma$ and any chain $c \in D_\xi$, c is eventually-contained in σ, so that each class of D_ξ corresponds to a unique class in D_σ. Thus, for example, the cardinality of $D_\xi/\sim$ is an invariant of Σ.

Definition 3.15 (codimension of an a.e. class) *Let $\Sigma \in \Re H$ be any almost-equality class of H°. The* codimension of Σ in H *is defined to be the cardinality $D_\xi/\sim$ for some (and hence any) $\xi \in \Sigma$.*

For example, the principal class is characterized by being of codimension zero.
Our current objective is to characterize the classes of finite codimension in terms of the order structure we have for $\Re H$.

Suppose now that $\xi \in \Sigma$, where $\Sigma \in \Re H$ is of finite non-zero codimension δ. Now, $D_\xi/\sim$ carries the domination partial ordering, and we may select an equivalence class by fixing a corresponding chain $c = \{c_n\}_{n=1}^{\infty}$. For each n, we let ξ_n be the result of truncating ξ at c_n.

We note the sequence $(\xi_n)_{n=1}^{\infty}$ converges in H° to the limit ξ, tempting us to wish for eventually-all of the ξ_n to belong to the same almost-equality class. Were this our case, then we would have found a $\Sigma_1 \in \Re H$ such that $\Sigma \cap \overline{\Sigma_1}$ is non-empty; by corollary 3.10, we would have $\Sigma_1 < \Sigma$ in $\Re H$.

Let us prove that this is indeed the case when one selects the equivalence class of c carefully enough: from now on, assume this class is maximal with respect to the domination ordering. The existence of such a class is guaranteed by our assumption that the codimension of Σ is finite.

Fix a complete set of representatives $c^{(1)} = c, \ldots, c^{(\delta)}$ in $D_\xi/\sim$. For each $i > 1$, we have $c^{(i)} \not\succeq c$, meaning that there exists $N_i \in \mathbb{N}$ such that no relation of the form $c_m^{(i)} \leq c_{N_i}$ is possible for any m. We choose N to be maximal among the N_i.

Given $n > N$, ξ_n may be obtained from ξ_{n+1} by truncating c_{n+1} (see lemma 3.13); thus, in order to show these two ultrafilters are almost-equal, it will be enough to prove that c_n does not lie on an infinite descending chain through ξ_{n+1}.

Suppose $d = \{d_m\}_{m=1}^{\infty}$ is a descending chain in ξ_{n+1} with the property that $d_1 = c_n$. Since d is a descending chain and H is discrete, its elements have to eventually leave the set $\{h \in H \,\big|\, c_{n+1}^* \leq h\}$, so that a terminal segment of d represents a chain-class in ξ. Now, d cannot be equivalent to c, and we conclude that $c_n = d_1 \geq c_m^{(i)}$ for some fixed i and any m large enough, contradicting the definition of N.

One last remark has to be made regarding this construction: by the construction of ξ_n (for $n > N$), it is evident that all the chains $c^{(i)}$ $(i > 2)$ survive in ξ_n, which means that by passing from Σ to Σ_1 we have reduced the codimension of Σ by precisely 1.

We summarize our last results in the following

Proposition 3.16 *Suppose $\Sigma \in \Re H$ has finite dimension $\delta > 0$. Then, there exists $\Sigma_1 \in \Re H$ such that:*

(*i*) $\Sigma_1 < \Sigma$, *and*

(*ii*) $codim(\Sigma_1) = \delta - 1$.

In particular, we have shown that Σ has codimension δ if and only if the shortest maximal descending chain in $\Re H$ starting at Σ (and terminating in Π) has length δ.

3.3.2 Natural median structure on $\Re H$.

We recall H° has the structure of a median algebra defined by the median operation

$$\mathbf{med}(\alpha, \beta, \gamma) = (\alpha \cap \beta) \cup (\alpha \cap \gamma) \cup (\beta \cap \gamma). \tag{35}$$

This operation is well-defined modulo almost-equality, inducing a median operation (satisfying all the required identities) on $\Re H$. It will be important to us that this operation is, in some sense, well-behaved with respect to the ordering we have defined on $\Re H$:

Lemma 3.17 (principal intervals) *Suppose $\Sigma \in \Re H$, and Π is the principal class in H°. Then*

$$[\Pi, \Sigma]_{med} = \left\{ \Sigma_1 \,\middle|\, \Sigma_1 \leq \Sigma \right\}, \tag{36}$$

where $[\cdot,\cdot]_{med}$ refers to the interval in $\Re H$ with respect to the median operation.

Proof : It follows from the properties of median functions that the interval $[\Pi, \Sigma]_{med}$ may be defined as (compare with the definition in the beginning of 2.2.4)

$$[\Pi, \Sigma]_{med} = \left\{ \mathbf{med}(\Pi, \Sigma, A) \,\middle|\, A \in \Re H \right\}. \tag{37}$$

Suppose now $A \in \Re H$, and consider $\Sigma_1 = \mathbf{med}(\Pi, \Sigma, A)$. We need to show two things: that $\Sigma_1 \leq \Sigma$, and that every $B \in \Re H$ with $B \leq \Sigma$ arises in the same manner as did Σ_1. To show that $\Sigma_1 \leq \Sigma$ it is enough to produce an element of Σ lying in the Tychonoff closure of Σ_1. Fixing $\alpha \in A$, $\sigma \in \Sigma$ we consider the ultrafilters

$$\sigma_\pi = \mathbf{med}(\pi, \sigma, \alpha) = (\alpha \cap \sigma) \cup (\alpha \cap \pi) \cup (\sigma \cap \pi) \tag{38}$$

in Σ_1 arising for different choices of $\pi \in \Pi$. Given a finite subset F of σ, use the denseness of Π in H° to find $\pi \in \Pi$ containing F. Then the formula for σ_π shows $F \subset \sigma_\pi$, proving σ lies in the Tychonoff closure of Σ_1, as desired.
Conversely, suppose $\Sigma_1 \leq \Sigma$, let $\Sigma_2 = \mathbf{med}(\Pi, \Sigma, \Sigma_1)$, and let us show $\Sigma_2 = \Sigma_1$. By what we have shown already, using the symmetries of the median operation, we see that both $\Sigma_2 \leq \Sigma$ and $\Sigma_2 \leq \Sigma_1$.
Therefore, it will be enough to prove $\Sigma_1 \leq \Sigma_2$ – that is, to show Σ_2 is contained in the Tychonoff closure of Σ_1. Consider an element $\sigma_2 \in \Sigma_2$ arising in the form

$$\sigma_2 = (\sigma_1 \cap \sigma) \cup (\sigma_1 \cap \pi) \cup (\sigma \cap \pi), \tag{39}$$

where $\pi \in \Pi$, $\sigma \in \Sigma$ and $\sigma_1 \in \Sigma_1$. If now F is a finite subset of σ_2 and $\pi \in \Pi$, as before, is a principal ultrafilter containing F, then our aim is to show that

$$\Sigma_1 \cap \bigcap_{f \in F} S_f \neq \varnothing. \tag{40}$$

Since F, being a subset of σ_2, is a filter base in H, our version of Helly's theorem (see corollary2.12) implies it is enough to show $\Sigma_1 \cap S_f$ is non-empty for every $f \in F$.
We finish the argument by way of contradiction. Suppose some $f \in F$ has $\Sigma_1 \subset S_f^*$. Since S_f^* is Tychonoff-closed, $\overline{\Sigma_1}$ – and therefore Σ, too – is contained in S_f^*. In particular, neither of the sets $\sigma_1 \cap \sigma$, $\sigma_1 \cap \pi$ and $\sigma \cap \pi$ contains f, implying σ_2 does not contain it – a contradiction. ■

We shall henceforth supress the subscript $_{med}$ in the notation for intervals in $\Re H$, knowing that both notions of an interval in $\Re H$ – both as a median algebra and as a partially ordered set – coincide.

Definition 3.18 (A notion of gcd in $\Re H$) *For $A, B, \Sigma \in \Re H$, we say Σ is a greatest common divisor for A and B, and denote this by $\Sigma = gcd(A, B)$, if*

1. $\Sigma \leq A$ *and* $\Sigma \leq B$, *and* –

2. *for any* $\Sigma' \in \Re H$ *satisfying* $\Sigma' \leq A$ *and* $\Sigma' \leq B$ *we have* $\Sigma' \leq \Sigma$.

Greatest common divisors indeed exist, as the following corollary of the last lemma shows:

Corollary 3.19 (existence of a gcd) *Any two elements* $A, B \in \Re H$ *have a gcd. Moreover,*

$$gcd(A, B) = \mathbf{med}(\Pi, A, B), \tag{41}$$

where Π *is the principal class of* H.

Proof : Set $\Sigma = \mathbf{med}(A, B, \Pi)$ and check it satisfies the required conditions.

1. by definition, $\Sigma \in [\Pi, A]$, and the last lemma implies $\Sigma \leq A$. The same argument works for B.

2. Since $\Sigma' \leq A, B$, we have, by the last lemma,

$$\begin{aligned} \Sigma' &\in [\Pi, A] \cap [\Pi, B] && (42)\\ &= [\Pi, \mathbf{med}(\Pi, A, B)] && (43) \end{aligned}$$

   (for the equality, see [Rol98], pg.15 claim **(Int 5)**), and use the same lemma again to deduce $\Sigma' \leq \Sigma$. ■

3.3.3 Restriction to a sub-poc-set.

Suppose K is a sub-poc-set of the ω-dimensional poc-set H (i.e., K is closed under $h \mapsto h^*$ and contains 0). We then have the restriction map res arising as the dual of the inclusion of K into H and preserving almost-equality. Consequently, we obtain a well-defined map $R^H_K : \Re H \to \Re K$.

Lemma 3.20 R^H_K *is a monotone non-decreasing map, i.e., if* $\Sigma_1 \leq \Sigma_2$ *in* $\Re H$, *then* $R^H_K \Sigma_1 \leq R^H_K \Sigma_2$ *in* K.

Proof : Saying $\Sigma_1 \leq \Sigma_2$ is equivalent to saying $\Sigma_2 \subset \overline{\Sigma_1}$. Since res is a continuous map, we conclude $res(\Sigma_2) \subseteq \overline{res(\Sigma_1)}$. Thus, the Tychonoff closure of $R^H_K \Sigma_1$ intersects $R^H_K(\Sigma_2)$, showing the desired inequality. ■

3.4 Convergence of principal ultrafliters.

Let us now consider some issues of convergence (in the Tychonoff topology) of sequences of (principal) ultrafilters.

Lemma 3.21 *Suppose* $(\sigma_n)_{n=1}^{\infty}$, $(\tau_n)_{n=1}^{\infty}$ *are convergent sequences in* H° *with limits* σ *and* τ, *respectively. If there exists* $M \in \mathbb{N}$ *such that* $\Delta(\sigma_n, \tau_n) \leq M$ *for all* n, *then* $\sigma \text{æ} \tau$.

Proof : Suppose $\sigma \smallsetminus \tau$ contains an infinite sequence $\{h_m\}_{m=1}^{\infty}$, and let us prove $\Delta(\sigma_n, \tau_n)$ is unbounded. Note that $h_m^* \in \tau$ for all m.
Given $m \in \mathbb{N}$ we define

$$B_m = \{h_1, \ldots, h_m\}\,, \tag{44}$$

producing a pair of disjoint open neighbourhoods $V(B_m)$ and $V(B_m^*)$ of σ and τ, respectively. By the definition of convergence in H°, there exists N such that $\sigma_n \in V(B_m)$ and $\tau_n \in V(B_m)$ hold for all $n \geq N$, showing $\Delta(\sigma_n, \tau_n)$ is unbounded. ■

The next result, shows how one can "optimize" convergent sequences of principal ultrafliters.

Lemma 3.22 (averaging lemma) *Suppose $(\pi_n)_{n=1}^{\infty}$ is an infinite sequence of principal ultrafilters converging to the limit $\sigma \in H^\circ$. Let Σ denote the almost-equality class of σ, and let pr_Σ denote the projection map of Π onto Σ.*
We define a new sequence of principal ultrafilters inductively as follows:

$$\begin{aligned}
\pi_1' &= \pi_1\,, && (45)\\
\pi_2' &= \mathbf{med}(\pi_1', \pi_2, \sigma)\,, && (46)\\
&\vdots && (47)\\
\pi_{n+1}' &= \mathbf{med}(\pi_n', \pi_{n+1}, \sigma)\,. && (48)
\end{aligned}$$

Then, the sequence we have defined converges on σ and satisfies the additional requirement that $\Delta(pr_\Sigma(\pi_n'), \sigma)$ is an eventually-zero monotone non-increasing function of n.

Proof : Suppose now A is a finite subset of σ. Find N such that $\pi_n \in V(A)$ for all $n \geq N$. But then, for all $n \geq N$ we must have $A \subseteq \pi_n \cap \sigma \subseteq \pi_n'$, showing $\pi_n' \in V(A)$, and proving that the π_n' do converge on σ.
Now let us compare $\pi_{n+1}' \vartriangle \sigma$ to $\pi_n' \vartriangle \sigma$:

$$\begin{aligned}
h \in \sigma \smallsetminus \pi_{n+1}' &\Leftrightarrow h \in \sigma \ \wedge\ h^* \in \pi_{n+1}' && (49)\\
&\Rightarrow h \in \sigma \ \wedge\ h^* \in \pi_n' \cap \pi_{n+1} && (50)\\
&\Rightarrow h \in \sigma \ \wedge\ h^* \in \pi_n' && (51)\\
&\Leftrightarrow h \in \sigma \smallsetminus \pi_n'\,. && (52)
\end{aligned}$$

Thus, $\sigma \smallsetminus \pi_{n+1}' \subseteq \sigma \smallsetminus \pi_n'$. Applying the complementation operator to both sides we obtain $\pi_{n+1}' \smallsetminus \sigma \subseteq \pi_n' \smallsetminus \sigma$, producing the inequality

$$\pi_{n+1}' \vartriangle \sigma \subseteq \pi_n' \vartriangle \sigma. \tag{53}$$

The result of applying pr_Σ to both sides of the inequality is obtained by intersecting both sides of the inequality with $T(\Sigma)$ (and then applying the natural isomorphism between $\overline{\Sigma}$ and $T(\Sigma)^\circ$), and we obtain:

$$pr_\Sigma(\pi_{n+1}') \vartriangle \sigma \subseteq pr_\Sigma(\pi_n') \vartriangle \sigma \tag{54}$$

This shows that the quantity $\Delta(pr_\Sigma(\pi'_n), \sigma)$ is, indeed, a non-increasing function of n taking values in $\mathbb{N} \cup \{0\}$. Moreover, the above inclusion shows that, if v is the eventual value of the latter quantity, then the above inequality implies there is an element $\sigma_1 \in \Sigma$ which equals eventually all the projections $pr_\Sigma(\pi'_n)$. Therefore, σ_1 is a limit of the sequence $pr_\Sigma(\pi'_n)$. Since pr_Σ is continuous and the π'_n converge to σ, the fact H° is Hausdorff implies $\sigma = \sigma_1$, showing $v = 0$, as desired. ■

We now consider a special case of convergence.

Lemma 3.23 (geodesic sequences converge) *Suppose $(\pi_n)_{n=1}^\infty$ is a geodesic ray in Π with respect to the metric d – that is, for all $m, n \in \mathbb{N}$ we have*

$$\Delta(\pi_n, \pi_m) = \big|n - m\big|. \tag{55}$$

Then $(\pi_n)_{n=1}^\infty$ converges to a limit $\sigma \in H^\circ$, and the averaged sequence $(\pi'_n)_{n=1}^\infty$ defined above coincides with $(\pi_n)_{n=1}^\infty$.

Proof : We may write $\pi_{n+1} = [\pi_n]_{a_n}$ for $a_n \in \min(\pi_n)$ and every $n \in \mathbb{N}$. The assumption that the π_n form a geodesic ray in Π means that

$$\pi_i \vartriangle \pi_j = \left\{a_j, \ldots, a_{i-1}, a_j^*, \ldots, a_{i-1}^*\right\} \tag{56}$$

for all $i > j$, with all the a_n distinct. In particular, for $j = 1, i = n > 1$ we have that the list

$$L_n := \pi_{n+1} \smallsetminus \pi_1 = \{a_1^*, \ldots, a_n^*\}, \tag{57}$$

consists of distinct n distinct elements and that π_m contains L_n for all $n > m$.
Let now $h \in H$ be a fixed halfspace. We shall show that either $h \in \pi_n$ for all but finitely many values of n, or $h^* \in \pi_n$ for eventually all n: indeed, if $h \notin \pi_n$ for all n, then $h^* \in \pi_n$ for all n and we are done; the same holds for h^*, so we may assume there are $i, j \in \mathbb{N}$ such that $h \in \pi_i$ and $h^* \in \pi_j$; without loss of generality, assume $i < j$, and we are forced to conclude that h^* appears on the list L_j, meaning h cannot appear on any of the lists L_n for $n > j$. Thus, $h^* \in \pi_n$ for all $n > j$, as desired.
We define an ultrafilter σ as follows: for any $h \in H$ we set $h \in \sigma$ if and only if $h \in \pi_n$ for all but finitely-many values of n. By what we have just shown, σ is, indeed, an ultrafilter on H, and since $(\pi_n)_{n=1}^\infty$ converges to σ pointwise, it converges to σ in the Tychonoff topology, as required.

Finally, let us use induction to prove $\pi'_n = \pi_n$ holds for all n, the base step ($n = 1$) having been taken care of by the definition of the "averaged" sequence. Given $n \in \mathbb{N}$, suppose that for all $k \leq n$ we have $\pi'_k = \pi_k$, and let us consider π'_{n+1}:

$$\begin{aligned}
\pi'_{n+1} &= (\pi'_n \cap \pi_{n+1}) \cup (\pi'_n \cap \sigma) \cup (\pi_{n+1} \cap \sigma) && (58)\\
&= (\pi_n \cap \pi_{n+1}) \cup (\pi_n \cap \sigma) \cup (\pi_{n+1} \cap \sigma) && (59)\\
&= (\pi_{n+1} \smallsetminus \{a_n^*\}) \cup (\pi_n \cap \sigma) \cup (\pi_{n+1} \cap \sigma). && (60)
\end{aligned}$$

Note that the second and third summands differ only by the element a_n^*, which, by the argument given above (set $h = a_n$) must lie in σ. Thus, π'_{n+1} contains π_{n+1} as a subset, implying they are equal, as desired. ■

Corollary 3.24 *Suppose $(\pi_n)_{n=1}^{\infty}$ is a geodesic ray in Π, and let σ be its limit. Further, let Σ be the almost-equality class of σ, and let pr_Σ denote the natural projection map of Π onto Σ. Then $pr_\Sigma(\pi_n)$ is eventually constant, and $\Delta(pr_\Sigma(\pi_n), \sigma)$ is monotone non-increasing and eventually-zero.* ■

The previous result has a converse. It is only natural to ask how far is an averaged sequence from being a geodesic vertex-path in Π:

Proposition 3.25 *Suppose $(\pi_n)_{n=1}^{\infty}$ is a sequence in Π converging to a limit π_∞ in H°, and let $(\pi'_n)_{n=1}^{\infty}$ denote the averaged sequence as defined above. Then $(\pi'_n)_{n=1}^{\infty}$ is a subsequence of a geodesic ray in Π.*

Proof : Given the elements π'_n, we select, for each $n \geq 1$, a geodesic vertex path p_n from π'_n to π'_{n+1}. Our guess is that the concatenation of all these paths is a geodesic ray in Π.
Let q_n denote the concatenation $p_1 * \ldots * p_n$. It will be enough to show that q_n is a geodesic path for all n. To do that, we use induction on n, where the case $n = 1$ is a sure success, by construction.
Suppose now that q_{n-1} is a geodesic vertex path in Π for some $n \geq 2$. In view of the induction hypothesis it is enough to prove the following equality:

$$\Delta(\pi'_1, \pi'_n) + \Delta(\pi'_n, \pi'_{n+1}) = \Delta(\pi'_1, \pi'_{n+1}).$$

Since q_{n-1} is a geodesic, to prove this equality it would be enough to show that

$$\Delta(\pi'_{n-1}, \pi'_n) + \Delta(\pi'_n, \pi'_{n+1}) = \Delta(\pi'_{n-1}, \pi'_{n+1}).$$

Applying the definition of d, this takes the form of:

$$\left|\pi'_{n-1} \vartriangle \pi'_n\right| + \left|\pi'_n \vartriangle \pi'_{n+1}\right| = \left|\pi'_{n-1} \vartriangle \pi'_{n+1}\right|.$$

Since $\pi'_{n-1} \vartriangle \pi'_{n+1} = (\pi'_{n-1} \vartriangle \pi'_n) \vartriangle (\pi'_n \vartriangle \pi'_{n+1})$, proving this would be the same as to prove that $\pi'_{n-1} \vartriangle \pi'_n$ and $\pi'_n \vartriangle \pi'_{n+1}$ do not intersect.
Note both sets are symmetric, and hence so is their intersection. Suppose now there was an element $h \in \pi'_{n+1} \smallsetminus \pi'_n$ lying in $\pi'_n \vartriangle \pi'_{n-1}$. Observe that then $h \in \pi'_{n-1} \smallsetminus \pi'_n$. Since

$$\pi'_{n+1} = (\pi'_n \cap \pi_{n+1}) \cup (\pi'_n \cap \pi_\infty) \cup (\pi_{n+1} \cap \pi_\infty),$$

and since $h \notin \pi'_n$, we conclude $h \in \pi_{n+1} \cap \pi_\infty$.
On the other hand, writing

$$\pi'_n = (\pi'_{n-1} \cap \pi_n) \cup (\pi'_{n-1} \cap \pi_\infty) \cup (\pi_n \cap \pi_\infty),$$

and observing $h^* \in \pi'_n \smallsetminus \pi'_{n-1}$ imply $h^* \in \pi_n \cap \pi_\infty$.
Thus, both h and h^* lie in π_∞ – a contradiction to π_∞ being an ultrafilter. ■

In order to complete the picture it remains to show a standard way of approximating a given element σ_∞ of a given class $\Sigma \in \Re H$ by an averaged sequence of elements of Π.

For every $\pi \in \Pi$, recall that $pr_\Sigma(\pi) = P(\Sigma) \cup (\pi \cap T(\Sigma))$ is an element of Σ. In fact, we know that $pr_\Sigma : \Pi \to \Sigma$ is a surjective median morphism, which allows to select $\pi \in pr_\Sigma^{-1}(\sigma_\infty)$.
Consider the intersection of $\min(\pi)$ with $P(\Sigma)^*$: for any $a, b \in \min(\pi)$, the only possible relations are $a \pitchfork b$ and $a^* \leq b$; however, if $\Sigma \subset S_{a^*}$, then $a^* \leq b$ would imply $\Sigma \subset S_b$, not allowing $\Sigma \subset S_{b^*}$. Thus, $\min(\pi) \cap P(\Sigma)^*$ is a transverse subset of $\min(\pi)$. We may then define a function

$$\left\{ \begin{array}{rcl} F_\Sigma : \Pi & \to & \Pi \\ \pi & \mapsto & [\pi]_{\min(\pi)\cap P(\Sigma)^*} \end{array} \right. \tag{61}$$

The sequence $\sigma_n = F_\Sigma^n(\pi)$ then necessarily converges to $pr_\Sigma(\pi)$ in the Tychonoff topology. Moreover, $(\sigma_n)_{n=1}^\infty$ is easily seen to be an averaged sequence once we notice that $\sigma_n \cap T(\Sigma)$ is constant and that the intersections $\sigma_n \cap P(\Sigma)$ form a strictly increasing sequence.

This construction is motivated by the notion of a normal cube path introduced by Niblo and Reeves in [NR98], with the σ_n constituting, in a sense, a normal cube path from π to $pr_\Sigma(\pi)$.

Definition 3.26 (canonical flow) *Given $\Sigma \in \Re\mathcal{H}$, the map $F_\Sigma : \Pi \to \Pi$ will be called the canonical flow on Π in the direction of Σ.*

The motivation for the name is, of course, the fact that iterating F_Σ over Π decomposes Π into orbits converging (through iterating F_Σ) onto the points of Σ.

4 Halfspace systems and boundary decompositions.

Suppose X is a proper $CAT(0)$-space.

Definition 4.1 *A* halfspace *$h \subset X$ is a non-empty open convex subset such that $h^* \triangleq X \smallsetminus \overline{h}$ is also convex. The intersection $\overline{h} \cap \overline{h^*}$ will be called* the wall associated with *h, and denoted by $W(h)$; if S is a set of halfspaces, then $W(S)$ will denote the set of walls $W(h)$ for $h \in S$. The sets $\varnothing, X$ are, by definition, the* trivial *halfspaces of X.*

Our notion of a halfspace system is the obvious generalization of what one observes in a cubing.

Definition 4.2 *A* halfspace system *in a CAT(0) space X is a family $\mathcal{H}$ of halfspaces containing the trivial halfspaces, ordered by containment, invariant under the operation $h \mapsto h^*$ and satisfying the following conditions:*

(**H1**) *Every point $x \in X$ has a neighbourhood U_x intersecting only finitely many walls associated with halfspaces of $\mathcal{H}$.*

(**H2**) *$\mathcal{H}$ contains no infinite transverse subfamily.*

The above notion of a halfspace system is discussed in detail in [Gur06]. The additional requirement $(H2)$ is required for us to be able to apply the results of the previous section.

For the sake of completeness, we provide all the definitions in this text, but we shall refer the reader there for the proofs of basic results.
The goal of this section is to explain how the Roller boundary $\Re\mathcal{H}$ of a halfspace system $\mathcal{H}$ induces a partition on the CAT(0)-boundary $\partial_\infty X$ of X, and to study some of the simplest connections between geometric properties of $\mathcal{H}$ and the behavior of ρ in order to establish a reasonably general framework within which to continue our discussion.
Two facts about geodesic rays and halfspaces in CAT(0) spaces are used throughout this work without reference:

Lemma 4.3 *Suppose h is a halfspace in X, and suppose $c : [0, \infty) \to X$ is a geodesic ray emanating from a point of h. If $c(\infty) \in \partial_\infty h$, then the image of c lies in the closure of h. Moreover, if $c(T) \in W(h)$ for some $T > 0$, then $c(t) \in W(h)$ for all $t \geq T$, and $c(\infty) \in \partial_\infty W(h)$.* ■

Lemma 4.4 *For any halfspace h in X we have $\partial_\infty W(h) = \partial_\infty h \cap \partial_\infty h^*$.* ■

4.1 Defining the boundary decomposition map.

Given a halfspace system $\mathcal{H}$ on X and a point $\xi \in \partial_\infty X$ we consider the set

$$\mathcal{H}^\xi = \left\{ \sigma \in \mathcal{H}^\circ \;\middle|\; \forall_{h \in \sigma}\, \xi \in \partial_\infty \overline{h} \right\}. \tag{62}$$

Given any $\sigma \in \mathcal{H}^\xi$, the set

$$T(\xi) = \left\{ h \in \mathcal{H} \;\middle|\; \xi \in \partial_\infty \overline{h} \smallsetminus \partial_\infty W(h) \right\} \tag{63}$$

is obviously contained in σ. One should think of $W(h)$ as being transverse to ξ in the sense that for any point $p \in h^*$ the geodesic ray γ emanating from p and having endpoint ξ will inevitably cross $W(h)$ – hence the notation for $T(\xi)$. An immediate observation following from the preceding lemmas is:

Lemma 4.5 *Let $\sigma \in \mathcal{H}^\circ$ and $\xi \in \partial_\infty X$. Then $\sigma \in \mathcal{H}^\xi$ if and only if σ contains $T(\xi)$.* ■

Proof : We had already noticed that $T(\xi) \subseteq \sigma$ for all $\sigma \in \mathcal{H}^\xi$. Suppose now that $\sigma \in \mathcal{H}^\circ$ contains $T(\xi)$, and let us show it lies in $\mathcal{H}^\xi$. Thus, taking $h \in \sigma \smallsetminus T(\xi)$ we must show $\xi \in \partial_\infty \overline{h}$. Since σ contains $T(\xi)$, we must have $h^* \notin T(\xi)$, and we have that neither of

$$\xi \in \partial_\infty \overline{h} \smallsetminus \partial_\infty W(h)\,, \qquad \xi \in \partial_\infty \overline{h^*} \smallsetminus \partial_\infty W(h) \tag{64}$$

holds. This implies $\xi \in \partial_\infty W(h)$, finishing the proof. ■

Corollary 4.6 *$\mathcal{H}^\xi$ is non-empty for all $\xi \in \partial_\infty X$.*

Proof : $T(\xi)$ is clearly a filter base, and is therefore contained in an ultrafilter. ■

Continuing the analysis of $\mathcal{H}^\xi$ let us define

$$P(\xi) = \{0, 0^*\} \cup \left\{ h \in \mathcal{H} \;\middle|\; \xi \in \partial_\infty W(h) \right\}. \tag{65}$$

Observe that $P(\xi)$ inherits the structure of a discrete ω-dimensional poc-set. By defintion, $\sigma \in \mathcal{H}^\xi$ implies $\sigma \subset T(\xi) \cup P(\xi)$. Thus, a halfspace $a \in \min(\sigma)$ will lie either in $T(\xi)$ or in $P(\xi)$:

- if $h \in \min(\sigma) \cap P(\xi)$, then $[\sigma]_h$ remains in $\mathcal{H}^\xi$;
- if $h \in \min(\sigma) \cap T(\xi)$, we inevitably have $[\sigma]_h \notin \mathcal{H}^\xi$.

The first observation hints at a possible connection between the structure of $\mathcal{H}^\xi$ and that of $P(\xi)^\circ$, realized through the map res_ξ dual to the inclusion $i : P(\xi) \hookrightarrow \mathcal{H}$. We shall now study the restriction of res_ξ to $\mathcal{H}^\xi$ more closely.
For any $h \in P(\xi)$ and $k \in T(\xi)$, we have that ξ is an interior point of k and a boundary point of h at the same time, which means $k < h$ is impossible; the relation $k < h^*$ is impossible for the same reason, as $h^* \in P(\xi)$. Thus, for all $h \in P(\xi)$ and $\sigma \in \mathcal{H}^\xi$ we have

$$h \in \min(\sigma) \Leftrightarrow h \in \min(\sigma \cap P(\xi)) = \min(res_\xi \sigma). \tag{66}$$

We conclude that for every almost-equality class $\Sigma \in \Re\mathcal{H}$ intersecting $\mathcal{H}^\xi$, the map res_ξ maps $\Sigma \cap \mathcal{H}^\xi$ *onto* the corresponding almost-equality class of $P(\xi)^\circ$.
Next, suppose $\bar{\sigma} \in P(\xi)^\circ$, and let us find a $\sigma \in \mathcal{H}^\xi$ satisfying $\bar{\sigma} = res_\xi(\sigma)$. We shall use our previous observation that no element of $T(\xi)$ is smaller than an element of $P(\xi)$. Define $\sigma = T(\xi) \cup \bar{\sigma}$, and observe it automatically satisfies (**UF1**). In order to check (**UF2**) it is enough to take $h \in \bar{\sigma}$ and $k \in T(\xi)$ and rule out $h < k^*$ and $k < h^*$: applying the involution we see that the two are equivalent, and the second contradicts the observation we have just made. Thus we have proved:

Proposition 4.7 *Suppose $\mathcal{H}$ is a halfspace system in a proper CAT(0)-space X, let $\xi \in \partial_\infty X$ be any point, and let $\Re(\xi)$ denote the set of almost-equality classes represented in $\mathcal{H}^\xi$. Then, res_ξ induces a one-to-one order-preserving correspondence R_ξ of $\Re(\xi)$ onto $\Re P(\xi)$, with res_ξ mapping $\Sigma \cap \mathcal{H}^\xi$ bijectively onto $R_\xi(\Sigma)$. In particular, $\Re(\xi)$ has a minimum element – the unique $\Sigma \in \Re(\xi)$ satisfying $R_\xi(\Sigma) = \Pi_{P(\xi)}$.* ■

Thus the last proposition enables us to define a map of $\partial_\infty X$ into the Roller boundary of $\mathcal{H}$:

Definition 4.8 (boundary decomposition map) *Suppose $\mathcal{H}$ is a halfspace system in a proper CAT(0)-space X. The* boundary decomposition map

$$\rho : \partial_\infty X \to \Re\mathcal{H} \tag{67}$$

is defined to be the map sending each $\xi \in \partial_\infty X$ to the minimum element of $\Re(\xi)$.

The following examples provide the basic feeling of how this map works.

Example 4.9 *Let $\mathcal{H}$ be the halfspace system arising from the standard cubing of Euclidean 2-space and denote by $x^\pm$ and $y^\pm$ the positive and negative boundary points of the $x-$ and $y-$axes, respectively. For any $\xi \in \partial_\infty \mathbb{E}^2$ other than $x^\pm, y^\pm$ and every $h \in \mathcal{H}$ we have either $h \in T(\xi)$ or $h^* \in T(\xi)$, showing $\mathcal{H}^\xi$ consists of a unique point, and one easily sees there are precisely four such points; on the other hand, for $\xi = x^+$, say, one has $h \in T(\xi)$ for all positive vertical $h \in \mathcal{H}$, whereas all horizontal elements of $\mathcal{H}$ lie in*

$P(\xi)$. Thus, H^ξ is homeomorphic to the ultrafilter space of the standard cubing of $\mathbb{R}$, and ρ maps the boundary point x^+ to the "positive horizontal" almost-equality class of $\mathcal{H}^\circ$.

Note that if $\mathcal{H}$ in the above example were taken to consist only of the horizontal halfspaces of the standard cubing of $\mathbb{E}^2$, we would have all points ξ of $\partial_\infty \mathbb{E}^2$ except $y^\pm$ accept the principal class as their value under ρ, while $y^\pm$ would be mapped to the corresponding (two) non-principal classes.

Example 4.10 *The same thing happens also in $X = \mathbb{H}^2$, when $\mathcal{H}$ is taken to be the family of halfspaces corresponding to the system of walls one obtains by translating a hyperbolic line ℓ by a hyperbolic translation g with axis orthogonal to ℓ. The line ℓ separates $g(+\infty)$ from $g(-\infty)$, making $\mathcal{H}$ into a poc-set isomorphic to the standard poc-set structure on $\mathbb{Z}$, but here, given a point $\xi \in \partial_\infty X \smallsetminus \{g(\pm\infty)$ one is not able to associate a geodesic ray γ satisfying $\gamma(\infty) = \xi$ with* every *ultrafilter of the class $\rho(\xi)$.*

Another interesting example is one of a 'parabolic' nature:

Example 4.11 *Again, let X be the hyperbolic plane, and let ℓ be a line with an endpoint p on the ideal boundary; denote by $h_\pm$ the different components of $X \smallsetminus \ell$. Further let g be a parabolic translation fixing p, and let $\mathcal{H}$ be the orbit of $\{h_\pm\}$ under $\langle g \rangle$. Then, for any ideal point $\xi \neq p$ we shall have $\rho(\xi)$ is the principal class Π of $\mathcal{H}^\circ$, with $\Re(\xi)$ containing only this class, while for $\xi = p$ the set $\Re(\xi)$ equals $\mathcal{H}^\circ$, resulting in $\rho(\xi) = \Pi$ once again. Thus, here the representation map does not distinguish among the points of $\partial_\infty X$, though the assignment $\xi \mapsto \Re(\xi)$ does.*

An immediate application of the construction of ρ explains some of the behaviour of the above examples.

Proposition 4.12 *A point $\xi \in \partial_\infty X$ satisfies $\rho(\xi) = \Pi$ if and only if $T(\xi)$ contains no infinite descending chains.*

Proof : If $\rho(\xi) = \Pi$, then there exists $\sigma \in \Pi \cap \mathcal{H}^\xi$. Since $T(\xi) \subseteq \sigma$, this implies $T(\xi)$ cannot contain an infinite descending chain. Conversely, if $\rho(\xi) > \Pi$, then an element $\sigma \in \rho(\xi) \cap \mathcal{H}^\xi$ has to contain an infinite descending chain c. However, this chain cannot contain a subchain of elements from $P(\xi)$, as $\rho(\xi) \cap \mathcal{H}^\xi$ is mapped by the restriction map into the principal class of $P(\xi)^*$. Since $\sigma \subset T(\xi) \cup P(\xi)$, all but finitely many elements of the chain c must lie in $T(\xi)$. ■

4.2 Uniform systems.

The motivation for our construction of ρ was provided by known examples of cubings as well as by examples of tilings of the plane derived from the obvious Coxeter groups. We will henceforth restrict our attention to halfspace systems possessing a certain geometric property that is automatically satisfied when X is a cubing and $\mathcal{H}$ is its natural halfspace system. The same property is also enjoyed by the halfspace system defined on the Davis-Moussong complex of an infinite Coxeter system of finite rank: take X to be the Davis-Moussong complex of a Coxeter system (W, R) of finite rank, and let $\mathcal{H}$ be the system of halfspaces arising as the set of complementary components of the walls. The

example we have chosen for illustrating most of the work done in this paper is that of the regular hexagonal tiling of the Euclidean plane $\mathbb{E}^2$, which is nothing else than the Davis-Moussong complex of the Coxeter system

$$W \cong \left\langle r, s, t \,\middle|\, r^2, s^2, t^2, (rs)^3, (rt)^3, (st)^3 \right\rangle \ , \tag{68}$$

as illustrated in figure 1. Walls are defined to be the fixed point sets of reflections of

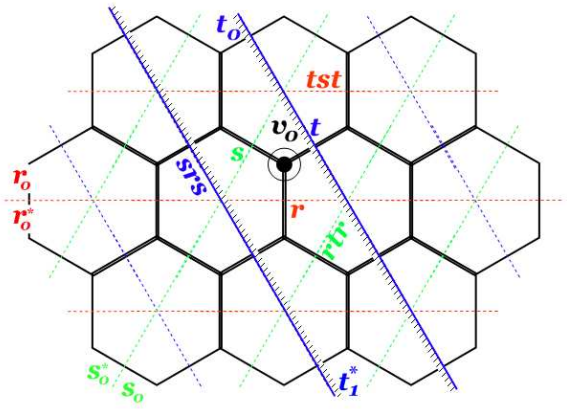

Figure 1: The hexagonal packing of $\mathbb{E}^2$: walls are the fixed-point sets of reflections; $\mathcal{H}$ decomposes as the union of three systems of proper halfspaces $\{r_n\}, \{s_n\}, \{t_n\}$ (and complements) indexed by $n \in \mathbb{Z}$, with $\{r_a, s_b, t_c\}$ transverse for any $a, b, c \in \mathbb{Z}$; we set r_0, s_0, t_0 to be three pairwise transverse minimal halfspaces among those containing the vertex v_0 corresponding to the unit element of W.

the system (W, R), and it can be shown (for example, see [Wil98]), that this system of walls coincides with $W(\mathcal{H})$.

The work of Brink and Howlett [BH93] shows that this particular class of examples has the *parallel walls property*, which, in the case of Coxeter groups was introduced by Davis and Shapiro [DS91]:

Definition 4.13 (parallel walls property) *A halfspace system $\mathcal{H}$ in a proper CAT(0) space X has the* parallel walls property*, if there exists a constant $C > 0$ such that for every $h \in \mathcal{H}$ and $x \in X$ satisfying $d(x, h^*) > C$ there exists a halfspace $k \in \mathcal{H}$ such that $x \in k < h$.*

Remark 4.14 *When we are mentioning parallel walls, this should not be mistaken for walls lying at a bounded Hausdorff distance from each other. By a pair of parallel walls we only mean walls $W(h), W(k)$ not being separated one from the other by another wall of $\mathcal{H}$.*

With respect to the visual boundary $\partial_\infty X$ of X, $\mathcal{H}$ having the parallel walls property becomes the following –

Definition 4.15 (conical points, uniformness) *A point $\xi \in \partial_\infty X$ is said to be a* conical limit point of $\mathcal{H}$, *$T(\xi)$ is non-empty, and for any $a \in T(\xi)$ and any cone neighbourhood U of ξ in $X \cup \partial_\infty X$ there exists $b \in T(\xi)$ satisfying $b < a$ and $b^* \cap U \neq \varnothing$.*
A halfspace system $\mathcal{H}$ on a proper CAT(0) space X is said to be uniform, *if all points of $\partial_\infty X$ are conical limit points of $\mathcal{H}$.*

Example 4.16 *A uniform halfspace system satisfies $\rho^{-1}(\Pi) = \varnothing$, by proposition 4.12.*

The following lemma in [Gur06] verifies that a halfspace system with the parallel walls property and satisfying $T(\xi) \neq \varnothing$ for all $\xi \in \partial_\infty X$ is uniform:

Lemma 4.17 *Suppose $\mathcal{H}$ is a halfspace system in a proper CAT(0) space X, having the parallel walls property, and let $\xi \in \partial_\infty X$. Then, for every $h \in T(\xi)$ and every cone neighbourhood U of ξ in $\hat{X}$ there exists a $k \in T(\xi)$ such that $k < h$ and $k^* \cap U \neq \varnothing$.* ■

Let us get back to the relation between our general situation and the situation in Coxeter groups. Given a halfspace system $\mathcal{H}$, we notice that most points of the space do not lie on any wall of $\mathcal{H}$. For any such point $x \in X$ it is possible to associate its *(closed) chamber –*

Definition 4.18 (chamber) *If $\mathcal{H}$ is a halfspace system in a CAT(0) space X, and $x \in X$ does not lie on any wall of $\mathcal{H}$, the chamber $ch(x)$ of x is defined as the intersection of closures of all halfspaces in $\mathcal{H}$ containing x.*

Example 4.19 *Clearly, since $\rho^{-1}(\Pi) = \varnothing$, a uniform halfspace system has bounded chambers.*

In the Davis-Moussong $X = M(W, R)$ complex of a Coxeter system (W, R), all chambers are bounded, as every chamber corresponds to a unique element of W, and W acts cocompactly on X. The following result from [Gur06] establishes the connection between uniformness and the special case we had just considered:

Theorem 4.20 *Let X be a proper CAT(0) space and let $\mathcal{H}$ be a halfspace system in X. If $\mathcal{H}$ has bounded chambers and satisfies the parallel walls property, then $\mathcal{H}$ is uniform. If, in addition, X admits a geometric group action by a group G stabilizing $\mathcal{H}$ and $\mathcal{H}$ is uniform, then $\mathcal{H}$ has the parallel walls property.* ■

4.2.1 Technical tools for uniform systems.

We now derive some criteria for comparing the images of points $\xi, \eta \in \partial_\infty X$ under the decomposition map ρ associated with a uniform system $\mathcal{H}$.

Lemma 4.21 *Suppose $\mathcal{H}$ is a halfspace system in a proper CAT(0) space X, and let $\xi, \eta \in \partial_\infty X$. Then,*

1. *If $T(\xi) \subseteq T(\eta)$, then $\rho(\xi) \leq \rho(\eta)$;*
2. *If $\mathcal{H}$ is uniform, then $\rho(\xi) \leq \rho(\eta)$ holds if and only if both $P(\eta) \subseteq P(\xi)$ and $T(\xi) \subseteq T(\eta)$.*

Proof : Suppose $T(\xi) \subseteq T(\eta)$. Then $\rho(\eta) \in \Re(\xi)$, and since $\rho(\xi)$ is the minimum of $\rho(\xi)$, we get $\rho(\xi) \leq \rho(\eta)$.
Now, if $\mathcal{H}$ is uniform, recall that for $\Sigma \in \Re\mathcal{H}$, $T(\Sigma)$ denotes the poc-set of all $h \in \mathcal{H}$ which are either trivial or satisfy $h \pitchfork \Sigma$ (i.e. $\Sigma \not\subset S_h$ and $\Sigma \not\subset S_h^*$). Consider $T(\rho(\xi))$: the uniformness of $\mathcal{H}$ implies $T(\rho(\xi)) = P(\xi)$ for all $\xi \in \partial_\infty X$.
Now assume $\rho(\xi) \leq \rho(\eta)$. Then $T(\rho(\xi))$ must contain $T(\rho(\eta))$, thus proving $P(\xi)$ contains $P(\eta)$. Next, recall $\rho(\xi)$ is contained in the closed set $\mathcal{H}^\xi = \bigcap_{h \in T(\xi)} S_h$; since $\rho(\eta)$ is contained in the closure of $\rho(\xi)$, every element of $\rho(\eta)$ contains $T(\xi)$. Thus, $T(\eta) \cup T(\xi)$ is a filter base in $\mathcal{H}$. Since $\mathcal{H} = T(\eta) \cup T(\eta)^* \cup P(\eta)$ holds together with $T(\xi) \cap P(\eta) = \varnothing$, we conclude $T(\xi) \subset T(\eta)$. ■

Corollary 4.22 *Suppose $\rho(\xi) \leq \rho(\eta)$. Then*

$$T(\xi) \subseteq T(\eta) \subseteq T(\xi) \cup P(\xi). \tag{69}$$

Proof : We have to show $h \in T(\eta)$ contradicts $h \in T(\xi)^*$. Take $\sigma \in \rho(\eta)$. Thus, S_h is a neighbourhood of σ and $\rho(\xi) \leq \rho(\eta)$ implies there exists $\tau \in \rho(\xi)$ lying in S_h, implying $h \in T(\xi)^*$ is impossible. ■

Lemma 4.23 *Suppose $\mathcal{H}$ is a uniform halfspace system in a proper CAT(0) space X, and let $\xi, \eta \in \partial_\infty X$. If $T(\eta) \subseteq T(\xi) \cup P(\xi)$ and $P(\eta) \subseteq P(\xi)$, then $\rho(\xi) \leq \rho(\eta)$.*

Proof : Let us show $\mathcal{H}^\eta \subset \mathcal{H}^\xi$, which will imply $\rho(\xi) \leq \rho(\eta)$. Consider the restriction map $res_\eta : \mathcal{H}^\eta \to P(\eta)^\circ$: we know it is an isomorphism of median algebras. If $\sigma \in \mathcal{H}^\eta$, consider $\sigma \cap P(\xi)$: being a filter base, this set is contained in an element $\bar{\sigma} \in P(\xi)^\circ$. Set $\tilde{\sigma} = \bar{\sigma} \cup T(\xi) \in \mathcal{H}^\xi$, and observe this is an ultrafilter containing all the elements of σ. Thus, $\sigma = \tilde{\sigma}$, and we have $\sigma \in \mathcal{H}^\xi$. ■

4.2.2 The closure formula.

Proposition 4.24 *Suppose $\mathcal{H}$ is a uniform halfspace system in a proper CAT(0) space X, and let $\Sigma \in \Re\mathcal{H}$ lie in the image of the boundary decomposition map ρ. Then –*

$$(\mathbf{FF_0}) \quad \overline{\rho^{-1}(\Sigma)} \subseteq \bigcup_{\Sigma_1 \leq \Sigma} \rho^{-1}(\Sigma_1). \tag{70}$$

In particular, since $\rho^{-1}(\Pi) = \varnothing$, then $\rho^{-1}(\Sigma)$ is closed for all $\Sigma \in \Re\mathcal{H}$ of codimension 1.

Proof : Let $\xi \in \overline{\rho^{-1}(\Sigma)}$, and denote $\Sigma_1 = \rho(\xi)$. In order to show $\Sigma_1 \leq \Sigma$, it will be enough to show that Σ is represented in $\mathcal{H}^\xi$: since Σ_1 is a minimum in $\Re(\xi)$, we will have $\Sigma_1 \leq \Sigma$. Thus, we must find $\sigma \in \Sigma$ containing $T(\xi)$.
Let $(\xi_n)_{n=1}^\infty$ be a sequence of boundary points satisfying $\rho(\xi_n) = \Sigma$ and converging to ξ. Recall that each of the $\mathcal{H}^{\xi_n}$ is a union of almost-equality classes, and so one may fix an element $\sigma \in \Sigma$ and deduce that $T(\xi_n) \subseteq \sigma$ for all n.
Now, if $h \in T(\xi)$, then there exists a neighbourhood U of ξ in $\partial_\infty X$ disjoint from $\partial_\infty \overline{h^*}$. Thus, $\xi_n \in U$ for n large enough implies $h \in T(\xi_n)$, proving $h \in \sigma$. ■

The above formula is, in a way, the main attribute of our decomposition map, demonstrating a relationship between the possibility of disconnecting large 'chunks' (fibers of ρ) of $\partial_\infty X$ from each other and the structure of the image of $\partial_\infty X$ under ρ as a partially-ordered set. We will see that more precision can be achieved for uniform systems.

4.2.3 restriction to a subspace

Many examples of CAT(0) spaces arise as a result of various gluings. This is the reason why it might be a good startegy to compute the decomposition of the boundary of such a space based on the knowledge of the boundary decompositions of the subspaces being glued together to obtain the whole space. Proposition 4.28 of this paragraph gives the precise tool for making such computations, showing that the boundary decomposition of any closed convex subspace of X is the one induced from the decomposition of X via the natural inclusion map.
Let us fix a closed convex subspace F of X, and consider the map r_F^X from $\mathcal{H}$ into the set of halfspaces of F defined by

$$r_F^X(h) = \begin{cases} h \cap F & ; h^* \cap F \neq \varnothing \\ F & ; h^* \cap F = \varnothing. \end{cases} \tag{71}$$

We shall use the symbol r to denote r_F^X whenever the choices of X and F are unambiguous.

Definition 4.25 *Suppose $\mathcal{H}$ is a halfspace system in the proper CAT(0) space X, and suppose F is a closed convex subspace of X. We then define the restriction $\mathcal{H}\big|_F$ of $\mathcal{H}$ to F as the set of all (possibly trivial) halfspaces of F of the form $r_F^X(h)$ defined above for $h \in \mathcal{H}$.*

Lemma 4.26 *Suppose $\mathcal{H}$ is a halfspace system in the proper CAT(0) space X, and F is a closed convex subspace of X. Then the restriction $\mathcal{H}\big|_F$ of $\mathcal{H}$ to F is indeed a halfspace system in the space F.*

Proof : Obviously, $\mathcal{H}\big|_F$ forms a poc-set under inclusion and the restricted complementation operator. Let us denote $h \cap F$ by $h\big|_F$ for all $h \in \mathcal{H}$, and let the image of any $S \subset \mathcal{H}$ under the map $h \mapsto h\big|_F$ be denoted by $S\big|_F$.
In order to make sure $\mathcal{H}\big|_F$ is a halfspace system in F, we need to verify conditions (**H1**), (**H2**). We note that $S\big|_F$ is nested whenever $S \subset \mathcal{H}$ is nested, and, accordingly, S is transverse whenver $S\big|_F$ is transverse. Thus, applying axiom (**H2**) for $\mathcal{H}$ we conclude there are no infinite transverse subsets in $\mathcal{H}_F$, showing $\mathcal{H}\big|_F$ is ω-dimensional. The condition (**H1**) clearly holds. ■

We now consider the dual r° of $r = r_F^X$. We note that r is finite-to-one, except, possibly, for the preimages of $0, 0^* \in \mathcal{H}\big|_F$ being infinite: indeed, only a finite number of distinct walls of $\mathcal{H}$ corresponds under r to any given non-trivial wall of $\mathcal{H}\big|_F$, because any family of halfspaces with intersecting walls in $\mathcal{H}$ is transverse. This implies r° maps almost-equality classes of $\mathcal{H}\big|_F^\circ$ into almost-equality classes of $\mathcal{H}^\circ$, since for any pair $\sigma, \sigma' \in \mathcal{H}\big|_F^\circ$,

if $\sigma \triangle \sigma'$ is finite then

$$r^{\circ}\sigma \triangle r^{\circ}\sigma' = r^{-1}(\sigma \triangle \sigma') \tag{72}$$

is a finite set, too (since $0^* \in \sigma \cap \sigma'$, all the fibers of r over $\sigma \triangle \sigma'$ are finite). Moreover, since r is a surjection (by definition), we may conclude that if $r^{\circ}\sigma \triangle r^{\circ}\sigma'$ is finite, then $\sigma \triangle \sigma'$ must be finite.
Thus, r° induces an injective map $\Re\mathcal{H}\big|_F \to \Re\mathcal{H}$. This map will be of some importance to us, so let us fix some notation for it:

Definition 4.27 *Suppose $\mathcal{H}$ is a halfspace system in the proper CAT(0) space X, F – a closed convex subspace of X, and let $i_F : \partial_\infty F \to \partial_\infty X$ denote the inclusion map. Then the map $\Re\mathcal{H}\big|_F \to \Re\mathcal{H}$ induced by r° will be denoted by $\Re(i_F)$.*

We use the map $\Re(i_F)$ to understand the relation between the Roller decomposition maps associated with F and X:

Proposition 4.28 *Suppose $\mathcal{H}$ is a halfspace system in the proper CAT(0) space X, and F is a closed convex subspace of X. Let ρ and ρ_F denote the Roller decomposition maps associated with X and F, respectively. Then, the following diagram is commutative:*

$$\begin{array}{ccc} \partial_\infty F & \xrightarrow{\rho_F} & \Re\mathcal{H}\big|_F \\ {\scriptstyle i_F}\big\downarrow & & \big\downarrow{\scriptstyle \Re(i_F)} \\ \partial_\infty X & \xrightarrow{\rho} & \Re\mathcal{H} \end{array} \tag{73}$$

Proof : We need to show $\Re(i_F) \circ \rho_F = \rho \circ i_F$. For this we fix a point $\xi \in \partial_\infty F$, and consider some $h \in \mathcal{H}$: obviously, if $h \cap F$ contains ξ in its boundary, then h contains ξ in its boundary; thus, if $\alpha \in \mathcal{H}\big|_F^{\xi}$, then $r^{\circ}\alpha = r^{-1}(\alpha)$ lies in $\mathcal{H}^{\xi}$, and we have

$$r^{\circ}(\mathcal{H}\big|_F^{\xi}) \subseteq \mathcal{H}^{\xi}, \tag{74}$$

implying $\rho(\xi) \leq (\Re(i_F) \circ \rho_F)(\xi)$.
To prove equality, assume, by way of contradiction, that $\Re(i_F)(\rho_F(\xi)) > \rho(\xi)$, and let $\alpha \in r^{\circ}(\rho_F(\xi))$. Since $\alpha \in \mathcal{H}^{\xi}$, $\alpha \subset T(\xi) \cup P(\xi)$, and there exists an infinite descending chain $(c_n)_{n=1}^{\infty}$ of elements of $P(\xi)$ in α (as opposed to any element of $\rho(\xi)$, which may not contain such a chain). If α_n is a sequence of elements in $\rho(\xi)$ converging on α, then we may assume α_n contains c_n for all n.
Consider the chain $\bar{c}_n = r(c_n)$. Either $c_n \cap F$ is a dense (and open) subset of F (possibly the whole of F) for all n, or $c_n \cap F$ and $c_n^* \cap F$ are both non-empty for n large enough. In the first case, taking a point $p \in F$ and a point $q \in c_1^*$ we find out that the geodesic segment $[p, q]$ in X intersects each of the walls $W(c_n)$, contradicting condition (**H1**) for $\mathcal{H}$. Thus, only the second case is possible, but then, taking $\bar{\alpha} \in \rho_F(\xi)$ such that $\alpha = r^{\circ}(\bar{\alpha})$, we find out that $(r(c_n))_{n=N}^{\infty}$ is a descending chain in $\bar{\alpha}$ for some N. This, too, is impossible, because no element of $\rho_F(\xi)$ may contain a descending chain of elements in $r(P(\xi))$, by the minimality of $\rho_F(\xi)$ (see proposition 4.7). ■

5 Connectivity properties of uniform halfspace systems.

From now on assume G is a group acting properly-discontinuously and co-compactly by isometries on X, keeping a uniform (and ω-dimensional!!) halfspace system $\mathcal{H}$ invariant.

5.1 Lines and flat sectors.

We recall a standard fact about CAT(0) spaces satisfying our assumptions:

Proposition 5.1 ([BH99] II proposition 9.5(3)) *For all $\xi, \xi' \in \partial_\infty X$ there exists a point $p \in X$ and representative rays $c \in \xi$ and $c' \in \xi'$ emanating from x_0 such that $\angle(\xi, \xi') = \angle_p(c, c')$. In particular, if $\angle(\xi, \xi') = \pi$, then there exists a geodesic line $c : \mathbb{R} \to X$ such that $\xi = c(\infty)$ and $\xi' = c(-\infty)$.*

Lemma 5.2 *Suppose $\rho^{-1}(\Pi) = \varnothing$, and suppose $c : \mathbb{R} \to X$ is a geodesic line with endpoints $\xi = c(\infty)$ and $-\xi = c(-\infty)$. Then $\rho(\xi)$ and $\rho(-\xi)$ are incomparable.*

Proof : Given the line c, we will show that $\rho(\xi) \leq \rho(-\xi)$ implies $\rho(\xi) = \Pi$, producing a contradiction.
For any $h \in T(\xi)$, note that if $c(\mathbb{R}) \nsubseteq \overline{h}$, then $-\xi$ cannot lie in $\partial_\infty W(h)$, implying $h^* \in T(-\xi)$, which is the same as $h \in T(-\xi)^*$ and implies $h \notin T(-\xi)$.
Now, the idea of the proof is to use the fact that $T(\xi)$ is non-empty to produce an element $h \in T(\xi) \smallsetminus T(-\xi)$, making the inequality $\rho(\xi) \leq \rho(-\xi)$ impossible: if this inequality holds, then corollary 4.21 implies $T(\xi)$ is contained in $T(-\xi)$, contradicting the existence of such an h.
In view of the remark above we may assume that every $h \in T(\xi)$ closure-contains the line $c(\mathbb{R})$. We may also assume no element of $T(\xi)$ lies in $P(-\xi)$, because this would also end the proof. Thus, we continue operating under the assumption that $c(\mathbb{R})$ is contained in h, for all $h \in T(\xi)$. But this is impossible: by uniformness, since $T(\xi)$ is non-empty there exists an infinite descending chain $(c_n)_{n=1}^{\infty}$ in $T(\xi)$; picking a point $p \in c_1^*$ and a point $q \in c(\mathbb{R})$, we immediately obtain a contradiction to (**H1**). ■

Corollary 5.3 (ρ separates π-discrete sets) *Suppose $\rho^{-1}(\Pi) = \varnothing$, and A is a subset of $\partial_\infty X$ which is π-discrete in the angular metric. Then the restriction of ρ to A is injective.* ■

The last corollary shows that the Roller boundary of a uniform halfspace system is particularly interesting when $\partial_\infty X$ is, say, connected with respect to the angular metric.

Example 5.4 *Suppose $X = \mathbb{H}^n$, and fix $x_0 \in X$. For any $x \in X$ and $g \in G$ let $h_{x,g}$ denote the halfspace of all points $p \in X$ which are closer to x than to gx. Then the set of halfspaces $\mathcal{H}$ consisting of all $h_{x,g}$ with $x \in G \cdot x_0$ and $g \in G$ is a uniform halfspace system in X satisfying $\rho^{-1}(\Pi) = \varnothing$. Since any two points of $\partial_\infty X$ are a distance π apart in the angular metric, we conclude ρ is one-to-one.*
Thus, in order for ρ to be meaningful for a group G acting on $\mathbb{H}^n$, one needs G to have parabolic points (and then $\mathcal{H}$ stops being uniform). Note that the same kind of problem will arise for any visibility space, whenever $\mathcal{H}$ is rich enough to satisfy $\rho^{-1}(\Pi) = \varnothing$.

5.2 Improved closure formula.

Let us now revisit the formula ($\mathbf{FF_0}$) we had derived in the previous section. It turns out that now we are able to make it into an equality.

Theorem 5.5 *Suppose G is a group acting properly and co-compactly by isometries on a proper CAT(0)-space X, and $\mathcal{H}$ is a G-invariant uniform halfspace system in X. Then, for all $\Sigma \in \rho(\partial_\infty X)$ one has the equality*

$$(\mathbf{FF}) \quad \overline{\rho^{-1}(\Sigma)} = \bigcup_{\Sigma_1 \leq \Sigma} \rho^{-1}(\Sigma_1)\,. \tag{75}$$

Proof : Suppose $\Sigma_1 < \Sigma$, and consider a boundary point ξ satisfying $\rho(\xi) = \Sigma_1$. Let us write $\Sigma = \rho(\eta)$ for some $\eta \in \partial_\infty X$. By corollary 4.21 and 4.23, we have

1. $P(\eta) \subset P(\xi)$, and –
2. $T(\xi) \subset T(\eta) \subset T(\xi) \cup P(\xi)$.

In particular, $T(\eta)$ is non empty.
Now, since $\rho(\eta)$ and $\rho(\xi)$ are comparable, lemma 5.2 implies we have $\angle(\xi, \eta) < \pi$. Since there exists in X a point p such that $\angle_p(\xi, \eta) = \angle(\xi, \eta)$, we conclude that the ideal triangle $\triangle\,(p, \xi, \eta)$ in X bounds a flat sector F in X (see [BH99], cor.9.9, p.283). For any boundary point ζ, let us denote the geodesic ray from p to ζ by $[p, \zeta]$.
Let us now consider any point $\zeta \in \partial_\infty F$ other than ξ. If we show that $\rho(\zeta) = \rho(\eta)$, then we are done, since ξ may clearly be written as the limit of a sequence of such ζ.
We fix $\zeta \in \partial_\infty F$ and note the ray $[p, \zeta]$ lies in the convex hull F of $[p, \xi]$ and $[p, \eta]$.
First, for any $h \in P(\eta) = P(\eta) \cap P(\xi)$ we note that since there exists $A > 0$ such that both $[p, \xi]$ and $[p, \eta]$ are eventually-contained in $N_A(W(h))$, so does $[p, \zeta]$, by the convexity of the metric in X; this proves $P(\eta) \subseteq P(\zeta)$. Next, if $h \in T(\eta)$, then from $h \in T(\xi) \cup P(\xi)$ we deduce in a similar fashion that h cannot lie in $T(\zeta)^*$, and we have shown $T(\eta) \subset T(\zeta) \cup P(\zeta)$, proving $\rho(\zeta) \leq \rho(\eta)$.

Let us now use $\rho(\xi) \leq \rho(\eta)$ more explicitly: there necessarily exists an infinite descending chain $(c_n)_{n=1}^\infty$ of elements in $P(\xi) \cap T(\eta)$. Selecting points $p_n \in [p, \eta] \cap c_n$ we necessarily have that $(p_n)_{n=1}^\infty$ converges on η – otherwise there would be a geodesic segment $[p, q] \subset [p, \eta) \cap X$ intersecting all the walls $W(c_n)$. Letting $\gamma_n = [p_n, \xi]$, we observe that $[p, \zeta]$ intersects each γ_n in a point $q_n \in F$, and that these points converge on ζ.
Now consider $h \in T(\xi)$. We already have $T(\xi) \subset T(\eta) \subseteq T(\zeta) \cup P(\zeta)$. Assume $h \in P(\zeta)$. Then, there exists $B > 0$ such that all but finitely many points of the sequence q_n belong to $N_B(h^*)$. By the convexity of the metric, all the rays γ_n are then eventually contained in $N_B(h^*)$, contradicting $h \in T(\xi)$. Thus we have shown $T(\xi) \subseteq T(\zeta)$ holds, which also implies $T(\zeta) \subseteq T(\xi) \cup P(\xi)$ and $P(\zeta) \subset P(\xi)$.

Finally we are in position to show $T(\eta) \subseteq T(\zeta)$, which will finish the proof. We already have $T(\eta) \subseteq T(\zeta) \cup P(\zeta)$, so let us assume there exists $T(\eta) \cap P(\zeta) \neq \varnothing$ and find out what goes wrong. We use the equality $P(\zeta) = T(\rho(\zeta))$ to deduce there exists an infinite descending chain of such elements in $T(\eta) \cap P(\zeta)$ – call it $(h_n)_{n=1}^\infty$. Without loss of

generality, $F \not\subseteq h_1$: if F were contained in all the h_n, then taking a point $p' \in h_1^*$, the geodesic segment $[p, p']$ would have intersected $W(h_n)$ for all n, which is impossible.
Since $h_n \in P(\zeta) \subseteq P(\xi)$, the intersection $F_n = F \cap \overline{h_n}$ is itself a flat sector in X having the same ideal boundary as F. Consider the ray $[p, \zeta]$: for large enough n we have $p \in h_n^*$ and $[p, \zeta]$ intersects $W(h_n)$ transversely; by the convexity of the metric, $h_n \in P(\zeta)$ is impossible, and we are done. ■

5.3 Boundary paths

Inspecting the above proof one immediately notes the connection between the structure of the Roller boundary of $\mathcal{H}$ and properties of "Euclidean" paths in the boundary. Adapting an idea of Croke and Kleiner ([CK00]), we define a special type of paths in the boundary of X as follows:

Definition 5.6 (safe paths) *A path $\alpha : [0,1] \to \partial_\infty X$ is said to be* safe *with respect to $\mathcal{H}$, if $\rho \circ \alpha([0,1])$ is a finite subset of $\Re\mathcal{H}$.*

Obviously, the concatenation of two safe paths is again safe, so that safe path-components of $\partial_\infty X$ are defined (with respect to a particular choice of $\mathcal{H}$).

Lemma 5.7 *safe path components of $\partial_\infty X$ are saturated with respect to ρ, i.e., for every safe path component C of $\partial_\infty X$ we have $\rho^{-1}(\rho(C)) = C$.*

Proof : It suffices to prove that for any $\xi, \eta \in \partial_\infty X$, if $\rho(\xi) = \rho(\eta)$, then ξ and η may be joined by a safe path. The procedure described in the proof of the theorem above gives the required path. ■

Proposition 5.8 (characterizing safe components) *Two points $\xi, \eta \in \partial_\infty X$ lie in the same safe path component if and only if there exists a sequence $\Sigma_0, \ldots, \Sigma_n$ of elements in $\rho(\partial_\infty X)$ satisfying*

1. *$\rho(\xi) = \Sigma_0$ and $\rho(\eta) = \Sigma_n$;*
2. *Σ_{i-1} and Σ_i are comparable for all $i \in \{1, \ldots, n\}$.*

We call such a sequence a connecting sequence of boundary classes of length n *for the points ξ and η.*

Proof : Once again, given the points ξ, η, the procedure used for the proof of (**FF**) shows that the existence of such a "connecting sequence" in $\rho(\partial_\infty X)$ implies there is a safe boundary path from ξ to η. For the converse, we claim that for any path $\alpha : [0,1] \to \partial_\infty X$, the set $\rho(\mathrm{Im}\alpha)$ – we shall abbreviate it as $\rho(\alpha)$ – contains a sequence of boundary classes connecting its endpoints.
Suppose not – then there exists a path providing a counter-example to our claim with $\rho(\alpha)$ of minimal possible size.
Let A be the set of all $i \in \{1, \ldots, n\}$ such that $\Sigma_j \leq \Sigma_i$ holds for no $j \neq i$. Then (**FF$_0$**) and the continuity of α imply the set $T_i = \alpha^{-1}(\rho^{-1}\Sigma_i)$ is a closed subset of $[0,1]$. Let J_i denote the closed subinterval of $[0,1]$ spanned by T_i.

Among all $i \in A$ find those with J_i maximal with respect to inclusion. on each such interval we use the previous lemma to redefine α in such a way that $\rho(\alpha(t)) = \Sigma_i$ for all $t \in J_i$. Thus, by the minimality property of α, all the J_i are disjoint.
For each $i \in A$, set $t(i) = \min(J_i)$, $t'(i) = \max(J_i)$, and let us write $A \smallsetminus \{0\} = \{i_1, \ldots, i_k\}$ with $t(i_1) < \ldots < t(i_k)$. Thus, we may break the interval $[0,1]$ into a series of consecutive intervals

$$\begin{array}{rclcrcl} I_0 & = & [0, t(i_1)], & & I_1' & = & [t(i_1), t'(i_1)] = J_{i_1} \\ I_1 & = & [t'(i_1), t(i_2)], & & I_2' & = & [t(i_2), t'(i_2)] = J_{i_2} \\ & \vdots & & & & & \\ I_k & = & [t'(i_k), 1], & & & & \end{array} \tag{76}$$

such that each interval $I'(\ell)$ is mapped by α into $\rho^{-1}(\Sigma_{i_\ell})$.
Suppose $k > 1$. Then for each of the intervals I_ℓ $(0 \leq \ell \leq k)$ the restriction α_ℓ of α to I_ℓ is a path satisfying $\rho(\alpha_\ell)$ is a proper subset of $\rho(\alpha)$, showing $\rho(\alpha_\ell)$ must contain a sequence of boundary classes connecting its endpoints. But then $\rho(\alpha)$ contains a sequence of boundary classes containing $\alpha(0)$ and $\alpha(1)$, which is impossible. Thus we deduce $k \geq 1$, implying $\rho(\alpha)$ has at most one minimal element – denote it by Σ_1 – except Σ_0 (which may, or may not be minimal).
Note now that $\Sigma_0 \leq \Sigma_n$ is impossible, by our assumption on α. Thus, Σ_1 does indeed exist (k cannot be zero).
Suppose now that Σ_1 is defined and Σ_0 is a minimal element of $\rho(\alpha)$. Then $[0,1] = I_0 \cup I_1' \cup I_1$ and $\rho(\alpha)$ may be written as the union of the two sets

$$B_0 = \left\{ \Sigma \in \rho(\alpha) \,\middle|\, \Sigma \geq \Sigma_0 \right\}, \quad B_1 = \left\{ \Sigma \in \rho(\alpha) \,\middle|\, \Sigma \geq \Sigma_1 \right\}, \tag{77}$$

and assume $\Sigma_n \notin B_0$ (else we are already done).
Now, ($\mathbf{FF_0}$) implies the sets B_0, B_1 are not disjoint, for otherwise $\rho^{-1}(B_0)$ and $\rho^{-1}(B_m)$ split $\mathrm{Im}\alpha$ as a union of two closed, non-empty disjoint sets, which is impossible.
Taking $\Sigma \in B_0 \cap B_1$, consider the path $\alpha' = \alpha\big|_{I_0 \cup I_1'}$: $\rho(\alpha')$ evidently contains the sequence $(\Sigma_0, \Sigma, \Sigma_1)$, which connects $\alpha(0)$ with $\alpha(t'(1))$. Next, the path $\alpha'' = \alpha\big|_{I_1}$ has $\rho(\alpha'')$ contain the sequence (Σ_1, Σ_n), which connects $\alpha(t'(1))$ with $\alpha(1)$. Thus, $\rho(\alpha)$ contains a sequence of boundary classes containing its endpoints – a contradiction.
Thus, we have shown $\rho(\alpha)$ has a unique minimal element $\Sigma_1 \neq \Sigma_0, \Sigma_n$, but then $(\Sigma_0, \Sigma_1, \Sigma_n)$ is a sequence connecting the endpoints of α, implying our α simply does not exist. ■

Thus, the safe path components of $\partial_\infty X$ correspond to comparability components in $\rho(\partial_\infty X)$, and the natural question to ask now is how many such comparability components are there in $\Re\mathcal{H}$.

5.3.1 Example: Flats.

Let X be as before, and suppose F is a flat – i.e., an isometric copy of $\mathbb{E}^m$ in X, for some positive m. We claim that $\partial_\infty F$ is safely path-connected. By lemma 4.26, $\mathcal{H}\big|_F$ is a halfspace system on F, and 4.28 tells us that $\rho_F(\partial_\infty F)$ is in one-to-one correspondence with the image of $\partial_\infty F$ under ρ in $\Re\mathcal{H}$. Thus it will be enough to show that $\partial_\infty \mathbb{E}^m$ is

safely path-connected with respect to any halfspace system on $\mathbb{E}^m$. Next, since geodesics in $\partial_\infty \mathbb{E}^m = SS^{m-1}$ are arcs of great circles, another application of the same two lemmas implies it is enough to show our claim for the case $m = 2$.
Suppose $\mathcal{H}$ is a halfspace system on $\mathbb{E}^2$, and for each $\alpha \in [0, \pi)$ denote by $\mathcal{H}_\alpha$ the subset of those $h \in \mathcal{H}$ such that the (positive) angle from the x-axis to $W(h)$ equals α. In $\mathbb{E}^2$ we necessarily have that, for all $\alpha, \beta \in [0, \pi)$ and all $h \in \mathcal{H}_\alpha, k \in \mathcal{H}_\beta$, $h \pitchfork k$ holds if and only if $\alpha \neq \beta$; since $\mathcal{H}$ contains no infinite transverse subset, we conclude $\mathcal{H}_\alpha$ is empty for all but finitely many values of $\alpha \in [0, \pi)$. Denote those values by $\{\alpha_1, \ldots, \alpha_d\}$, ordered in increasing order. With each α_i we associate two boundary points – denote them by $\pm\xi_i$, – the two endpoints of a line intersecting the x-axis at a (positive) angle α_i.
Now, for any point $\xi \in \partial_\infty \mathbb{E}^2$, if ℓ is a line with endpoint ξ, then for any $\alpha \in [0, \pi)$ and any $h \in \mathcal{H}_\alpha$ we have $h \in P(\xi)$ if and only if the angle of ℓ to the x-axis (measured from the positive ray of the x-axis to ℓ) equals α. Thus, $P(\xi)$ is non-trivial if and only if $\xi \in \{\pm\xi_i\}_{i=1,\ldots,d}$. Moreover, as ξ moves along the interior of a boundary arc whose endpoints are a pair of consecutive points of the cyclically-ordered set $\{\pm\xi_i\}_{i=1,\ldots,d}$, $T(\xi)$ remains constant, showing that $\rho(\xi)$ also remains constant. Thus, the image of $\partial_\infty \mathbb{E}^2 = SS^1$ under ρ consists of at most $4d$ elements, showing $\partial_\infty \mathbb{E}^2$ is safely path-connected. We have proved:

Proposition 5.9 (Euclidean boundaries are safe) *Suppose X is a proper CAT(0) space and $\mathcal{H}$ is a halfspace system on X. Then, for any isometrically embedded flat $F \cong \mathbb{E}^m$ in X, its boundary $\partial_\infty F$ in $\partial_\infty X$ is safely-connected with respect to $\mathcal{H}$, with $\rho(\partial_\infty F)$ isomorphic to $\Re(\mathbb{E}^M_\square)$ for some $M \geq m$, where $\mathbb{E}^M$ is taken with the standard halfspace system.* ■

Remark 5.10 *We have just proved that the Roller boundary of a halfspace system in a two-dimensional flat is finite. Using normal vectors (rather than angles) it is possible to generalize the same argument to prove that the Roller boundary of a halfspace system in a flat of any finite dimension is finite.*

In the case when X and $\mathcal{H}$ are invariant under a proper co-compact action by a group G of isometries of X, one way of obtaining a flat plane F is to find a subgroup A of G isomorphic to $\mathbb{Z}^2$ and then applying the Flat Torus Theorem to produce a flat F which is invariant under A. In this case one easily sees that the invariance of $\mathcal{H}$ under A implies the invariance of all the $\mathcal{H}_{\alpha_i}$ under A, showing $\mathcal{H}$ is poc-set isomorphic to the standard cubulation of $\mathbb{E}^d$. It is easy to see that the Roller boundary of this cubulation is isomorphic to the barycentric subdivision of a d-dimensional cube, and the image of $\partial_\infty F$ under ρ then traverses the cycle of length $4d$ spanned by the classes $\rho(\pm\xi_i)$.

5.3.2 Example: Flat sectors and uniform systems.

The preceding example explains the role of flats for a general halfspaces system, showing that boundaries of flats are always safe. That argument does not work for flat *sectors*, however, but it can be mended under the additional assumption that $\mathcal{H}$ satisfies the parallel walls property – in particular, when $\mathcal{H}$ is a uniform system invariant under a geometric group action.

Observe that for every $h \in \mathcal{H}$ for which the restriction to F is proper, the wall $W(h)$ intersects F in an interval separating F. The tool that makes the parallel walls property relevant is the following:

Lemma 5.11 *Suppose F is a flat sector and $h \in \mathcal{H}$ restricts to a proper halfspace of F. Suppose that γ is a geodesic ray in F such that $d(\gamma(t), F \cap \overline{h^*})$ is unbounded. Then $d(\gamma(t), h^*)$ is unbounded.*

Proof : Let v be the vertex point of F. If $d(\gamma(t), h^*)$ were bounded, that would imply $h \in P(\gamma(\infty))$. Let $x \in W(h) \cap F$ be any point. Since $W(h)$ is convex and complete, we must have that $[x, \gamma(\infty))$ is contained in $W(h)$. However, $x \in F$ and $\gamma \subset F$; since F is convex and complete, this implies $[x, \gamma(\infty))$ lies in F. Thus, $[x, \gamma(\infty))$ lies in $W(h) \cap F$, contradicting the assumption regarding γ. ■

Proposition 5.12 *Suppose X is a proper $CAT(0)$ space and $\mathcal{H}$ is a halfspace system in X, satisfying the parallel walls property. If F is a flat sector in X, then $\partial_\infty F$ is safely-connected with respect to $\mathcal{H}$.*

Proof : It is enough to prove the proposition for the case when F has a vertex angle $0 < \theta < \pi$. We will identify F with the set

$$\left\{ re^{i\alpha} \in \mathbb{C} \,\middle|\, r \geq 0\,, \quad 0 \leq \alpha \leq \theta \right\}\,,$$

with 0 corresponding to the vertex of F in X. To each angle in the interval $[0, \theta]$ corresponds a unique point $\xi_\alpha \in \partial_\infty F$. Same as before, we decompose (the proper part of) $\mathcal{H}\big|_F$ as the union of subsystems $\mathcal{H}_\alpha$ (throw in the trivial halfspaces) where $\alpha \in [0, \pi]$ and $h \in \mathcal{H}_\alpha$ if and only if the angle between the interval $W(h)$ and the x-axis (of F) equals α. Then it is clear that $\mathcal{H}_\beta \subseteq T(\xi_\alpha)$ whenever $\beta \neq \alpha$, and that $H_\alpha = P(\xi_\alpha)$. Thus, in order for $\partial_\infty F$ to be safely connected it is sufficient that $\mathcal{H}_\alpha$ be trivial for all but finitely many values of $\alpha \in (0, \theta)$.
For every $\alpha \in (0, \pi)$ and $h \in \mathcal{H}_\alpha$ we will denote the point of intersection of $W(h)$ with the boundary rays of F by $h(0)$, and the endpoint of the ray arising as $W(h) \cap F$ will be denoted by $h(\infty)$.
The assumption that $\mathcal{H}$ has no infinite transverse subset is used as follows. Suppose $h_n \in \mathcal{H}_{\alpha_n}$ $(n \in \mathbb{N})$ satisfy

1. $h_n(0) \in \mathbb{R}$ for all n,
2. $h_{n+1}(0) > h_n(0)$ for all n,
3. $\alpha_{n+1} > \alpha_n$ for all n.

Then it is clear that the rays $W(h_n) \cap F$ cross pairwise. Thus, such a configuration is impossible. A symmetric configuration (for which the rays $W(h_n) \cap F$ have $h_n(0) = r_n e^{i\theta}$, and with the α_n *decreasing*) is, of course, also impossible for the same reasons.
Consider the set A of all $\alpha \in (0, \theta)$ for which $\mathcal{H}_\alpha$ is poc-isomorphic to the standard halfspace system on $\mathbb{E}^1$. If A is infinite, extract a strictly monotone sequence α_n from A; because of the discreteness assumption on $\mathcal{H}$, it is then easy to construct (inductively) a forbidden configuration in $\mathcal{H}\big|_F$, producing a contradiction.

Therefore, it will be enough to prove that $\mathcal{H}_\alpha$ is poc-isomorphic to the standard halfspace system on $\mathbb{E}^1$ whenever it is non-trivial. From symmetry considerations it will be enough to prove that every $a \in \mathcal{H}_\alpha$ eventually-containing the ray of (positive) reals has a $b \in \mathcal{H}_\alpha$ satisfying $b < a$.
Fix a as above, and consider the ray of points of F we had identified with $\mathbb{R}_+$. By the preceding lemma, a point on this ray sufficiently far away from $W(a) \cap F$ must be contained in an element $h_1 \in \mathcal{H}$ satisfying $h_1 < a$. This implies that h_1 restricts to a proper halfspace of $\mathcal{H}\big|_F$, so that $h_1 \in \mathcal{H}_{\alpha_1'}$ for $\alpha_1' \leq \alpha$. We may take h_1 to be maximal with this property (as the interval $[h_1, a]$ is finite). In particular, there are no intermediate halfspaces between a and h_1.
If $\alpha_1' = \alpha$, we are done. If not, Then we apply the preceding lemma again, to the halfspace a and the ray $W(h_1) \cap F$. This results in a halfspace $h_2 < a$, $h_2 \in \mathcal{H}_{\alpha_2}$ and since a and h_1 had no intermediate halfspaces, we also obtain $\alpha_2 > \alpha_1$ and $h_2(0) > h_1(0)$. Proceeding inductively in the same manner, we obtain a forbidden configuration, as desired. ■

As a corollary, we obtain the following result:

Theorem 5.13 (Tits components and safe components coincide) *Suppose G is a group acting geometrically on a CAT(0) space X, and suppose $\mathcal{H}$ is a G-invariant uniform halfspace system. Then, for every $\xi \in \partial_\infty X$, the open Tits ball of radius π about ξ is contained in the safe component of ξ. In particular, the components of the Tits boundary of X coincide with the safe components of $\partial_\infty X$.* ■

Thus, the decomposition map ρ defined by $\mathcal{H}$ introduces a new structure on the Tits boundary, determined by the ordering of the Roller boundary associated with $\mathcal{H}$.

5.3.3 Example: $\mathbb{H}^2 \times \mathbb{H}^2$

Let X_1, X_2 be visible proper CAT(0) spaces, and let $X = X_1 \times X_2$. We claim:

Proposition 5.14 *If $\mathcal{H}$ is* any *uniform halfspace system on X, then $\partial_\infty X$ is safely-connected with respect to $\mathcal{H}$.*

Proof : the boundary $\partial_\infty X$ is naturally homeomorphic to the spherical join $\partial_\infty X_1 * \partial_\infty X_2$. Let $\xi = [\xi_1, \xi_2, \alpha]$ and $\eta = [\eta_1, \eta_2, \beta]$ be two distinct points of $\partial_\infty X$ (with their standard representations as points of $\partial_\infty X_1 * \partial_\infty X_2$, where $\xi_i, \eta_i \in \partial_\infty X_i$ and $\alpha, \beta \in [0, \pi/2]$).
Now, there are geodesic lines $\ell_i \subset X_i$ joining ξ_i to η_i, and we may consider the embedded flat $F = \ell_1 \times \ell_2 \subset X$: we note that $\partial_\infty F$ is then the join of the two-point subspaces $\{\xi_1, \eta_1\}$ and $\{\xi_2, \eta_2\}$. In particular, $\partial_\infty F$ contains our points ξ and η. By proposition 5.9, there is a safe path in $\partial_\infty X$ from ξ to η. ■

We would like to draw the attention of the reader to the fact that, in this example, the Roller boundary of $\mathcal{H}$ may be considered as an extremely highly-connected graph when both boundaries $\partial_\infty X_1$ and $\partial_\infty X_2$ are infinite. For example, given any finite set of points A in $\partial_\infty X$, and any pair of points $\xi, \eta \in \partial_\infty X \smallsetminus A$, there will be a safe path from ξ to η missing A, and this is independent of the choice of $\mathcal{H}$ (so long as $\mathcal{H}$ is uniform). On

the other hand, $\Re\mathcal{H}$ is by no means a trivial structure: If X and $\mathcal{H}$ admit a geometric action by some group, then since the equatorial copies of the $\partial_\infty X_i$ in $\partial_\infty X$ are infinite π-discrete sets, ρ is injective on each of them, producing pairwise-incomparable classes; this shows that, as a graph, $\rho(\partial_\infty X)$ may be quite complicated (an infinite complete bipartite graph in this case).

5.3.4 Example: the Croke-Kleiner examples

We recall the set of examples by Croke and Kleiner ([CK00]).
Given $\alpha \in (0, \pi/2]$, let T_1, T_2 be two standard ("square") geometric 2-tori (with the standard CW decomposition), and let T_0 be the geometric 2-torus obtained from a flat rhombus R with an interior angle α via the standard gluing; the two simple closed curves arising as the image of ∂R in T_0 will be denoted by m_1' and m_2'. Selecting meridional curves m_1, m_2 of unit length in T_1, T_2 respectively, we glue each T_i ($i \in \{1, 2\}$) to T_0 by identifying m_i with m_i' using an isometry. Let now X_α denote the universal cover of the resulting "torus complex" Y_α, and use Bridson's theorem to metrize X as a piecewise-Euclidean cell complex. It is known that X_α is a CAT(0) space, and we have that $\pi_1(Y_\alpha)$ which is independent of the choice of α) acts properly discontinuously and co-compactly by isometries on X_α. We let $\mathcal{H}$ be the halfspace system arising naturally from the cube structure on X_α, and it is evident that $\mathcal{H}$ is uniform, and suppress the index α until it is needed.
In their paper [CK00], Croke and Kleiner define a system $\mathcal{W}$ of *walls* in X – the distinct (disjoint) lifts of T_0 –, and a system $\mathcal{B}$ of *blocks* – all lifts of subsets of Y of the form $T_i \cup_{m_i = m_i'} T_0$. Let us color the 1-skeleton of X as follows: edges of X projecting to the curves $m_i = m_i'$ will be colored by 0, while edges projecting to curves not contained in T_0 will be colored according to their projections being contained in T_1 or T_2. The 1-skeleton of each block B is then colored using two colors (one of which is necessarily 0), according to the way it was constructed.
We shall now list some facts from [CK00].

local properties of blocks. Each block B is isometric to the cartesian product of a 4-regular metric tree – all of whose edges have length 1 – with $\mathbb{R}$, which is also standardly realized as a metric tree, whose vertex set equals $\mathbb{Z}$.

The fibers of the projection to the tree factor are called *singular fibers*. Note they are all geodesic lines in X, all parallel to each other.

global properties of blocks. The interaction among distinct blocks is as follows:

1. $\mathcal{B}$ covers X;
2. Two distinct blocks are either disjoint or share precisely one wall, which separates them in X;
3. The nerve of the covering $\mathcal{B}$ of X is a tree (actually, the Bass-Serre tree of the splitting of $\pi_1(X)$ as an amalgam over $\pi_1(T_0)$). Put more simply, the graph whose vertex set is $\mathcal{B}$ and whose edge set is $\mathcal{W}$, with an edge W incident to an edge B iff $W \subset B$, is a tree. We shall need the separation order on this tree, and write $a * b * c$ to denote that an element b of this tree separates the element a from the element c.

Using these facts, we may now proceed to give a precise description of $\mathcal{H}$. Given the structure of an individual block B, if a halfspace $h \in \mathcal{H}$ is such that both h and h^* intersect B, then either $h \cap B$ has the form of the cartesian product of a halfspace of $\mathbb{R}$ with the 4-regular tree corresponding to B, or $h \cap B$ is the product of a halfspace in the 4-regular tree with the whole of $\mathbb{R}$. In the former case, let us write $h \in T(B)$, while in the latter we shall write $h \in P(B)$, according to h being transverse or parallel to the singular fiber of the block B. $P(B)$, in turn, splits as the disjoint union of two sets $P_c(B)$ and $P_0(B)$, depending on whether $W(h)$ projects (under the covering $X \to Y$) into a non-zero (colored) edge, or into a zero-colored edge. Suppose now that B is a block. We consider the three cases discussed above for $h \in \mathcal{H}$ with respect to B, assuming neither h nor h^* contain B. The discussion is based on the intersection patterns among the liftings of the tori T_1, T_2 and T_0, as pictured in figure 2.

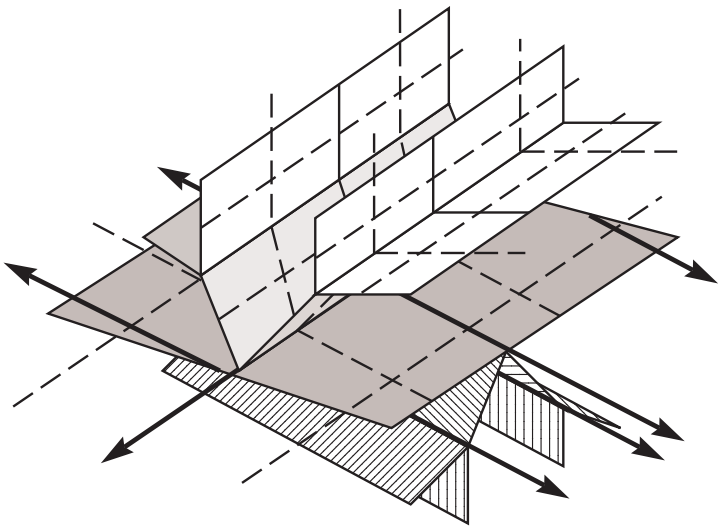

Figure 2: Intersection patterns of halfspaces and blocks in the Croke-Kleiner examples: the horizontal plane is the intersection plane of two neighbouring blocks (above and below the plane, respectively); singular fibers are denoted by arrows; walls (of the cubing) are denoted by dashed lines.

- $h \in T(B)$.
  In this case, for any block B' adjacent to B, we must have $h \in P_0(B')$, and $W = B \cap B'$ is the only wall of B' intersecting $W(h)$, because the intersection of $W(h)$ with W is a singular fiber of B'. In particular, for every $B'' \in \mathcal{B}$ satisfying $B * B' * B''$ we must either have $B'' \subset h$ or $B'' \subset h^*$.

  Thus, $T(B)$ induces non-trivial halfspace systems only on B and its immediate neighbours.

- $h \in P_c(B)$.

In this case, $W(h)$ is equal to a singular fiber of B. If $B' \neq B$ is any block and W is the unique wall of B satisfying $B * W * B'$, then we must have either $B' \subset h$ or $B' \subset h^*$ according to which of $W \subset h, W \subset h^*$ (resp.) occurs.

- $h \in P_0(B)$.

  This case is actually symmetric to the first one: there is a unique wall W of B containing the singular fiber $W(h) \cap B$, and therefore $h \in T(B')$, where B' is the unique block satisfying $B \cap B' = W$.

From this analysis we immediately deduce that to every proper element $h \in \mathcal{H}$ there corresponds a unique block $B \in \mathcal{B}$ satisfying either $h \in T(B)$ or $h \in P_c(B)$. Let us denote this block by $B(h)$, and decompose $\mathcal{H}$ as the disjoint union of subfamilies

$$\mathcal{H}_B = \left\{ h \in \mathcal{H} \,\middle|\, B(h) = B \right\} \cup \{\varnothing, X\}. \tag{78}$$

We note that $\mathcal{H}_B \cap \mathcal{H}_{B'}$, where B, B' is a pair of adjacent blocks, produces a halfspace system $\mathcal{H}_W$ (corresponding to the wall $W = B \cap B'$) such that the restriction map $r^X_W : \mathcal{H}_W \to \mathcal{H}\big|_W$ is an isomorphism, since for every $h \notin \mathcal{H}_W$ we have either $W \subset h$ or $W \subset h^*$.
We need now to compute $\mathcal{H}^\circ$ and the map ρ. Suppose $\alpha \in \mathcal{H}^\circ$ is not principal, so that α contains an infinite descending chain $(h_n)_{n=1}^\infty$; denote $B_n = B(h_n)$, and consider the following cases:

The sequence $(B_n)_{n=1}^\infty$ does not stabilize. Passing to an (equivalent) subsequence we may assume that B_n and B_{n+1} are disjoint for all n. From the analysis above it follows that we may find, for each n, a block B'_n satisfying $B_n * B'_n * B_{n+1}$ and a halfspace $h'_n \in P_c(B'_n)$ satisfying $h_{n+1} < h'_n < h_n$.

The sequence $(h'_n)_{n=1}^\infty$, apart from being equivalent to the original sequence $(h_n)_{n=1}^\infty$, also has the property that every $h \in HH$ must satisfy either $h'_n < h$ for sufficiently large n, or $h'_n < h^*$ for sufficiently large n. Consequently, the ultrafilter α is uniquely determined by the sequence, and, having no minimal elements, constitutes its own almost-equality class in $\mathcal{H}^\circ$. It follows that the class of α in $\Re\mathcal{H}$ is maximal, and that, in particular, no class arising in the same manner may be smaller. We will presently show that the class α determines in $\Re\mathcal{H}$ is, in fact, of codimension 1, implying it is its own comparability component in $\rho(\partial_\infty X)$, and that its preimage under ρ is a closed subset of $\partial_\infty X$. We shall call α a *singular* point of $\mathcal{H}^\circ$.

The sequence $(B_n)_{n=1}^\infty$ is eventually-constant. In this case, denote the terminal value of the sequence by B, and recall every block B is a closed convex subspace of X. By the definition of equivalence of chains, we may assume $B_n = B$ for all n. This allows us to consider the restriction $r^X_B : \mathcal{H} \to \mathcal{H}\big|_B$ together with the injective dual map of $\left(\mathcal{H}\big|_B\right)^\circ$ into $\mathcal{H}^\circ$. This map is a closed continuous embedding mapping almost-equality classes onto almost-equality classes, and we deduce that the almost-equality class of α cannot accumulate at a singular point of $\mathcal{H}^\circ$. This proves our previous assertion that the almost-equality class of a singular point of $\mathcal{H}^\circ$ is of codimension 1.

Now, since every element of $T(B)$ is transverse to every element of $P_c(B)$, we conclude that either $h_n \in T(B)$ for all n or $h_n \in P_c(B)$ for all n. Since both $T(B)$ and $P_c(B)$ are the proper elements of poc-sets associated to trees, an ultrafilter on $\mathcal{H}$ may not contain a pair of inequivalent chains of either type. Thus, the class corresponding to α in $\Re\mathcal{H}$ is necessarily of codimension at most 2, and we note that $\partial_\infty B$ is safely-connected with respect to $\mathcal{H}\big|_B$. This implies that the safe path-components of any two block-boundaries are equal, showing that, except for the safe components corresponding to the ends of the Bass-Serre tree, there exists only one more safe path-component in $\partial_\infty X$.

To conclude this example, we see that the map ρ does not distinguish among the boundaries of the spaces X_α, though it retains essential information regarding the structure of all these spaces.

5.4 Generalizing the examples

5.4.1 Surjectivity of ρ.

In all the preceding examples we have witnessed situations where ρ was essentially surjective: all non-principal classes were lying in the image of ρ. The following examples shows one cannot expect this to be true in general:

Example 5.15 (ρ is not necessarily surjective) *Let X be the subspace*

$$X = \bigcup_{n\in\mathbb{N}} [n-1, n] \times [n^2, \infty), \tag{79}$$

with the standard cubulation $\mathcal{H}$ induced from $\mathbb{E}^2$. There are three non-principal classes in $\Re\mathcal{H}$. However, this space is not almost-extendible, and has only one boundary point corresponding to vertical rays – both for the same reason: for any $C > 0$, if x is any point at a distance greater than C to the vertical ray emanating from the point $x_0 = (0, 1) \in X$, then there is no geodesic ray $[x_0, \xi)$ passing through $B(x, C)$.

Note that by a theorem of Ontaneda [Ont05], a CAT(0) space admitting a geometric action by a group is almost extendible. This is why we are tempted to ask the following question:

Question 5.1 *Suppose X is a proper cubing with standard halfspace system $\mathcal{H}$. If X is almost extendible, is it true that then every non-principal class in $\Re\mathcal{H}$ lies in the image of the boundary representation map?*

The results we have to report in this direction are much weaker, and serve, rather, as indications of situations in which one has an easy positive answer:

Proposition 5.16 *Suppose X is a proper cubing with standard halfspace system $\mathcal{H}$, and let $\Pi \neq \Sigma \in \Re\mathcal{H}$ be a class of codimension 1. Then $\Sigma \in \rho(\partial_\infty X)$.*

Theorem 5.17 *Suppose X is a proper cubing with standard halfspace system $\mathcal{H}$, and let $\Pi \neq \Sigma \in \Re\mathcal{H}$. If, for some $\pi \in \Pi$ the canonical flow $(F^n(\pi))_{n=1}^\infty$ fellow-travels a geodesic ray $[\pi, \xi)$ in X, then $\Sigma = \rho(\xi)$.*

Proof of prop 5.16: Let $\Sigma \neq \Pi$ be an element of $\Re\mathcal{H}$ and fix $\sigma_\infty \in \Sigma$. Let $(\sigma_n)_{n=1}^\infty$ be a sequence of ultrafilters in Π converging on σ_∞. Since X is proper, considering the σ_n as vertices of X we may pass to a subsequence such that the σ_n converge to a point $\xi \in \partial_\infty X$. As a result, $\sigma_\infty \in \mathcal{H}^\xi$, and we have

$$\Pi < \rho(\xi) \leq \Sigma\,. \tag{80}$$

In particular, $\rho(\xi) = \Sigma$ whenever $codim(\Sigma) = 1$. ■

Proof of theorem 5.17: We begin the proof in almost the same way as the previous one: let $\Sigma \neq \Pi$ be an element of $\Re\mathcal{H}$, let $F = F_\Sigma$ denote the canonical flow of Σ (defined on Π – see definition 3.26), and let $\pi \in \Pi$ be an ultrafilter satisfying $F^n(\pi) \in N_R([\pi, \xi))$ for all n for some $\xi \in \partial_\infty X$. Let $\sigma_\infty = pr_\Sigma(\pi) = \lim_{n\to\infty} F^n(\pi)$, where the limit is taken in $\mathcal{H}^\circ$ w.r.t. the Tychonoff topology.
As before, one must have $\rho(\xi) \leq \Sigma$. Suppose it were true that $\rho(\xi) < \Sigma$; then $P(\Sigma) \cap P(\xi)$ would have contained an infinite descending chain $(c_n)_{n=1}^\infty$. Without loss of generality we have $c_1^* \in \pi$, which, together with $c_1 \in P(\xi)$ implies $[\pi, \xi)$ is (closure-)contained in c_1^*. However, for any $n > R$ we have $d(c_{n+1}, c_1^*) > R$ (in (X, d)), which contradicts $F^{n+1}(\pi)$ lying in the $R-$neighbourhood of $[\pi, \xi)$. ■

Remark 5.18 *Of course, the assumption of the last theorem need not regard approximating sequences (for Σ) generated by the corresponding canonical flow. However, the examples we have indicate that this is the correct formulation (consider the preceding example where ρ is not surjective). The canonical flow seems to hold much information about the cubing X; for example, in any case when $\Delta(F_\Sigma(\pi), h) > \Delta(\pi, h)$ for all $h \in P(\Sigma)^*$ and $\pi \in \Pi$ (this is the case in $\mathbb{E}^n$ with the standard cubulation), it is easy to show that the sequence $F^n(\pi)$ lies on a geodesic. In a way, this is our motivation for hoping that canonical flows do indeed fellow-travel geodesic rays – at least for the case when X admits a proper co-compact cellular action by a group.*

5.4.2 Giant components.

The second feature common to the preceding examples was that the spaces and the cubings involved were all one-ended and the image of ρ in all cases contained a unique 'giant component' –

Definition 5.19 (giant component) *Suppose $\mathcal{H}$ is a halfspace system on a proper CAT(0) space X. A comparability component K of $\rho(\partial_\infty X)$ will be called a* giant component*, if its preimage in $\partial_\infty X$ under ρ is dense with respect to the cone topology.*

Remark 5.20 *As an example consider a situation when $\rho(\partial_\infty X)$ has a* finite *giant component K (like in the case of halfspace systems in $\mathbb{E}^d$): in this situation, the closure formula implies $K = \rho(\partial_\infty X)$. The example of the standard presentation 2-complex for $F_2 \times \mathbb{Z}$ shows that $K = \rho(\partial_\infty X)$ is also possible for infinite giant components.*

Theorem 5.21 *Suppose $\mathcal{H}$ is a uniform halfspace system on a proper CAT(0) space X. If $\rho(\partial_\infty X)$ has a giant component, then X is one-ended.*

Remark 5.22 *The converse is not true: for example, we have seen that for $X = \mathbb{H}^2$, any uniform halfspace system $\mathcal{H}$ makes any two values of ρ incomparable. Thus, every point in $\partial_\infty X$ corresponds to its own comparability component in spite of X being one-ended.*

Proof : Let K be a comparability component of $\rho(\partial_\infty X)$, and let $Y = \overline{\rho^{-1}(K)}$, the closure taken with respect to the cone topology. We show that any two points of $\xi, \eta \in Y$ belong to the same end of X: taking U and V to be cone neighbourhoods of ξ and η respectively, both not intersecting a ball $B(x_0, r) \subset X$ of a given radius $r > 0$, we find points $\xi' \in \rho^{-1}(K) \cap U$ and $\eta' \in \rho^{-1}(K) \cap V$. For some $R > r$ it is then possible to connect $\xi'(R)$ with $\eta'(R)$ with a piecewise-circular path. This path can be then extended on both ends by rays asymptotic to $[x_0, \xi)$ and $[x_0, \eta)$, which, by properties of rays in CAT(0) spaces, will necessarily be contained in U and V, respectively.
In particular, if K is a giant component, X is one-ended. ■

Theorem 5.23 (uniqueness of giant components) *Suppose $\mathcal{H}$ is a uniform halfspace system in a proper CAT(0) space X, and suppose F is a closed convex subspace of X such that*

1. *F coarsely separates X, and*
2. *$\rho_F(\partial_\infty F)$ has a unique comparability component.*

Then $\rho(\partial_\infty X)$ has at most one giant component.

Proof : The first assumption on F means that $\partial_\infty F$ separates $\partial_\infty X$. It will be enough to show that, if K is a giant component of $\rho(\partial_\infty X)$ then $\rho(\partial_\infty F)$ is contained in K. Since different comparability components of $\rho(\partial_\infty X)$ do not intersect, it will follow that K is the only giant component in stock.
Now, the second assumption on $\rho(\partial_\infty F)$, together with proposition 4.28, imply it is enough to show that $\rho(\partial_\infty F)$ intersects K.
Now, since F coarsely separates X, $\partial_\infty F$ coarsely separates $\partial_\infty X$. Since $\partial_\infty F$ is closed in $\partial_\infty X$, we may write $\partial_\infty X \smallsetminus \partial_\infty F = U \cup V$, where U and V are disjoint open subsets of $\partial_\infty X$. Since K is a giant component, there are points

$$\alpha \in \rho^{-1}(K) \cap U\,, \quad \beta \in \rho^{-1}(K) \cap V \tag{81}$$

and $\partial_\infty F$ intersects any (safe) path from α to β in $\partial_\infty X$. In particular, $\rho(\partial_\infty F)$ intersects K, and we are done. ■

Corollary 5.24 *Suppose a group G acts geometrically on a CAT(0) space, and suppose X admits a uniform G-invariant halfspace system $\mathcal{H}$. If G has a codimension-one free-abelian subgroup A of rank at least 2, then $\rho(\partial_\infty X)$ has at most one giant component, which, when it exists, is characterized as the comparability component containing the ρ-images of the boundaries of all coarsely-separating flats in X.*

Proof : By the flat torus theorem, there exists a flat F of dimension $r \geq 2$ in X, on which A acts co-compactly by translations; since A is a codimension-one subgroup, F coarsely separates X. Proposition 5.9 then implies that $\partial_\infty F$ is safely connected, allowing to apply the last theorem and its proof. ■

6 Co-compact actions.

Suppose, once again, that G is a group acting geometrically on the CAT(0) space X, and X has a G-invariant uniform halfspace system $\mathcal{H}$. In this section our aim will be to use the decomposition mapping $\rho : \partial_\infty X \to \Re\mathcal{H}$ for establishing when G acts co-compactly on the cubing dual to $\mathcal{H}$.

Theorem 6.1 *Suppose G is a group acting geometrically on a CAT(0) space X, and suppose $\mathcal{H}$ is a G-invariant uniform halfspace system. If every non-principal class of $\Re\mathcal{H}$ lies in the image of the boundary decomposition map $\rho : \partial_\infty X \to \Re\mathcal{H}$, then the action of G on $C(\mathcal{H})$ is co-compact.*

Note that the converse to this result is, in fact, a harder version of the problem discussed in the preceding section (question 5.1, theorem 5.17).

The first step in our discussion of the problem will be to reduce it to a combinatorial problem, which is done in a way analogous to that used by Williams in his thesis [Wil98]. Most of the technical details of this reduction, as well as some of the required results, were dealt with by the author in [Gur06]. The most important of these results states that $C(\mathcal{H})$ is locally finite.

In the following paragraph we provide an overview of the results and notions we shall use from that source.

6.1 Consistent ultrafilters.

For any point $x \in X$ let us consider the set B_x of all $h \in \mathcal{H}$ satisfying $x \in h$. B_x is a filter base. We define:

Definition 6.2 *A point $x \in X$ is said to* support *a subset $A \subseteq \mathcal{H}$, if $x \in \overline{h}$ for all $h \in A$. A subset $A \subset \mathcal{H}$ is said to be* consistent*, if it has a supporting point. We shall adopt the convention that the empty set is inconsistent.*

For ultrafilters, consistency is an easy matter:

Lemma 6.3 (consistent ultrafilters) *Suppose $\mathcal{H}$ is a halfspace system in a CAT(0) space X. Then:*

1. *any consistent ultrafilter is principal;*
2. *any point $x \in X$ supports an ultrafilter.*

The set of all consistent ultrafilters will be henceforth denoted by Π_0. ■

We see that the consistent ultrafilters – or, at least, those which are of the form B_x for some $x \in X$ – correspond to chambers. This, in effect, is what provides us with the combinatorial reduction of the co-compactness problem.

Lemma 6.4 *The set Π_0 is G-finite. In particular, there exists a constant $D > 0$, such that if $\pi, \pi' \in \Pi_0$ are supported on the same point, then $|\pi \triangle \pi'| \leq D$.* ■

Recall the metric Δ on Π. It is G-invariant. Now, since G acts on X stabilizing Π_0, it also stabilizes the level sets Π_n ($n \in \mathbb{N} \cup \{0\}$) of the function $\Delta(-, \Pi_0) : \sigma \mapsto \Delta(\sigma, \Pi_0)$. Since the action of G on Π_0 is co-finite, the action of G on Π will be co-bounded if and only if the function $\Delta(-, \Pi_0)$ is bounded. Since $C(\mathcal{H})$ is locally finite, its quotient by G will be compact if and only if it is bounded. Therefore, studying the growth of $\Delta(-, \Pi_0)$ will be our tool for studying the co-compactness problem.

At the base of the technique lies a geometric interpretation of $\Delta(\pi, \Pi_0)$:

Lemma 6.5 *Suppose $\pi \in \Pi$. Then $\Delta(\pi, \Pi_0) \leq n$ if and only if there exists a subset A of π of size n such that $\pi \smallsetminus A$ is consistent.* ■

Then, one needs a tool for understanding how $\Delta(\pi, \Pi_0)$ changes as one performs walks on the 1-skeleton of Π. For this we define –

Definition 6.6 *Suppose $\pi \in \Pi$ and $a \in \min(\pi)$. Denote*

$$\begin{aligned} a \in \min(\pi)_+ &\quad \Leftrightarrow \quad \Delta([\pi]_a, \Pi_0) > \Delta(\pi, \Pi_0)\,, \\ a \in \min(\pi)_- &\quad \Leftrightarrow \quad \Delta([\pi]_a, \Pi_0) < \Delta(\pi, \Pi_0)\,. \end{aligned}$$

Note that $\pi \in \Pi_0$ iff $\min(\pi)_-$ is empty. ■

The basic observations regarding the signed minimal sets of $\pi \in \Pi$ are:

Lemma 6.7 *For all $\pi \in \Pi$, the set $\min(\pi)_-$ is inconsistent.* ■

Corollary 6.8 (three ways to go down) *If $\pi \notin \Pi_0$, then $\min(\pi)_-$ contains at least three distinct elements.* ■

An important means of assessing distances to Π_0 is the growth of the following objects (as π recedes from Π_0):

Definition 6.9 (Shadows) *For all $\pi \in \Pi$, let the* shadow of π *be defined as the set*

$$\mathrm{sh}\,(\pi) = \left\{ \sigma \in \Pi_0 \,\middle|\, \delta(\sigma, \pi) = \Delta(\sigma, \Pi_0) \right\}, \tag{82}$$

and let the dual shadow $\mathrm{sh}^\circ\,(\pi)$ *of π be defined to be*

$$\mathrm{sh}^\circ\,(\pi) = \left\{ h \in \mathcal{H} \,\middle|\, \mathrm{sh}\,(\pi) \subseteq S_h \right\}. \tag{83}$$

Observe that $\mathrm{sh}^\circ\,(\pi)$ is a filter-base, and it is natural to expect that $\mathrm{sh}^\circ\,(\pi)$ be contained in π. If that is the case, it would mean that, as the distance of π from Π_0 increases, $\mathrm{sh}^\circ\,(\pi)$ diminishes accordingly, testifying to the growth of $\mathrm{sh}\,(\pi)$. This is exactly what we will require for the study of the co-compactness problem.

In order to gain a feeling of why shadows should be related to the co-compactness problem, let us consider once again the example of a Coxeter group (W, R) acting on its Davis-Moussong complex (and the associated halfspace system – call it $\mathcal{H}$ – which is, in fact, isomorphic to the corresponding system, ordered by the domination order

introduced in [BH93] for positive roots, and extended to all roots by [NR03]). It is now known, through the work of [Wil98] and [Cap05], that the action of W on $C(\mathcal{H})$ is co-compact if and only if W does not contain a Euclidean triangle subgroup. Let us consider such a subgroup W', and its associated Davis-Moussong (sub)complex, which is isometric to a Euclidean plane P tessellated by hexagons (corresponding to the finite dihedral parabolic subgroups): one notices that an inconsistent ultrafilter π will necessarily have $\min(\pi)_-$ define (through intersection) a triangle in P, and that these triangles can grow arbitrarily large; the fact that these triangles come in infinitely many different sizes shows that the number of conjugacy classes (in W', and hence in W) of subgroups $W'' < W$ isomorphic to W' must be infinite (as W acts on X by isometries) – which hints at Wilson's co-compactness criterion (W acts co-compactly on $C(\mathcal{H})$ iff W contains only finitely many conjugacy classes of triangular subgroup of each admissible type). Wilson and Caprace went on to study the structure of relations among reflections creating this abundance of 'similar shapes' (The hyperbolic triangle groups, lack this abundance of similar shapes, for the obvious geometric reasons; in fact, Caprace's work exhausts the possibilities for such 'shapes' in the general case, showing that the only reason for non-co-compactness is the presence of infinite similarity classes of Euclidean 'shapes'). In our, more general case, however, we shall have to study the 'shapes' themselves – these are modelled by shadows.

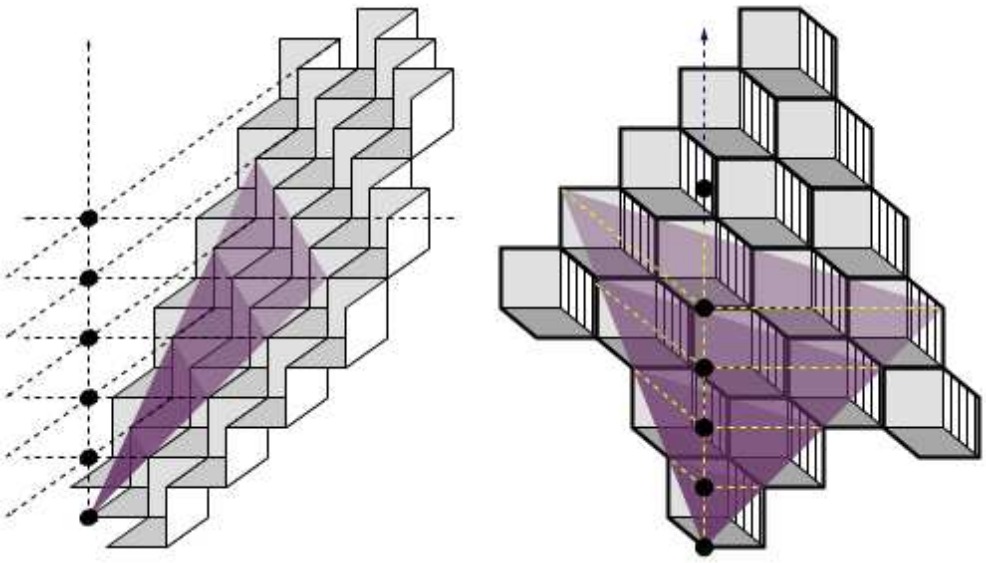

Figure 3: The hexagonal packing in $\mathbb{E}^2$ is the Davis-Moussong complex of a Coxeter triangle group for which $C(\mathcal{H})$ is isomorphic to the standard cubulation of $\mathbb{E}^3$. Each vertex v of the packing corresponds to the ultrafilter of halfspaces containing v; so that the Cayley graph of W (which is the 1-skeleton of the packing) is embedded in $C(\mathcal{H}) \cong \mathbb{E}^3$ as shown above (thick lines, right view), while inconsistent ultrafilters give rise to shadows.

The technicalities we require are summarized in a result we had proved in [Gur06]:

Lemma 6.10 *For every* $\pi \in \Pi$ *one has* $\min(\pi)_+ \subseteq \mathrm{sh}^\circ(\pi) \subseteq \pi$. *Furthermore, for any* $a \in \min(\pi)$ *one has:*

1. *if* $a \in \min(\pi)_+$ *then* $\mathrm{sh}(\pi) \subsetneq \mathrm{sh}([\pi]_a)$ *and* $\mathrm{sh}^\circ([\pi]_a) \subsetneq \mathrm{sh}^\circ(\pi)$*;*
2. *if* $a \in \min(\pi)_-$ *then* $\mathrm{sh}^\circ(\pi) \cap \min(\pi)_- = \varnothing$.

6.2 Proof of theorem 6.1

6.2.1 Limits of geodesic rays.

We need a geometric lemma involving the structure of $T(\xi)$ in a uniform halfspace system.

Lemma 6.11 *Suppose* $A \subseteq T(\xi)$ *is a non-empty set containing elements* $a_1, \ldots, a_m$ *such that* $A \smallsetminus \{a_1, \ldots, a_m\}$ *is consistent. Then* A *is a consistent set.*

Proof : It is enough to prove the result for the case $m = 1$. Suppose A were inconsistent. This would mean that

$$Y := \bigcap_{a \in A \smallsetminus \{a_1\}} \overline{a} \subset a_1^*. \tag{84}$$

Now, on one hand, Y is consistent, closed and convex, and so it contains a geodesic ray $[y, \xi)$, which, therefore lies entirely in a_1^*. On the other hand we see that $\overline{a_1}$ contains a ray $[x, \xi)$, implying $a_1 \in P(\xi)$ – contradiction. ■

The following proposition is the crucial ingredient in the proof of theorem 6.1. In a way, it states which geodesic vertex paths in the one-skeleton of the cubing dual to $\mathcal{H}$ produce limits representing elements of $\Re\mathcal{H}$ which do not lie in the image of ρ. The motivation for this proposition was given by the example of the hexagonal packing, as shown in figure 4.

Proposition 6.12 (escaping vertex paths) *Suppose* Π *contains a geodesic sequence* $(\pi_n)_{n=0}^\infty$ *satisfying*

$$\Delta(\pi_{n+1}, \Pi_0) > \Delta(\pi_n, \Pi_0)$$

for all n*, and let* Σ *denote the almost-equality class of the limit* π_∞ *of this ray. Then* $\Sigma \notin \rho(\partial_\infty X)$.

Proof : Suppose, on the contrary, that $\Sigma = \rho(\xi)$ for some $\xi \in \partial_\infty X$.
By lemma 3.23 and corollary 3.24, truncating an initial segment of the sequence $(\pi_n)_{n=0}^\infty$ we may assume that $\pi_n \cap P(\xi)$ is constant (and then, for all n, we would have $\pi_n \cap P(\xi) = \pi_\infty \cap P(\xi)$, which is a principal ultrafilter on $P(\xi)$), implying that all the elements of $\mathcal{H}$ whose orientations are being reversed along this ray of Π are halfspaces belonging to $T(\xi)^*$. Let us denote these halfspaces by a_n, so that we have $\pi_{n+1} = [\pi_n]_{a_n}$ and $a_n^* \in \min(\pi_{n+1})_-$ for all n.
Now observe that, for any n and $h \in \pi_n$, if $h \notin \mathrm{sh}^\circ(\pi_n)$, then $h \in \pi_{n+1}$ (otherwise, this would imply that $h \in \min(\pi_n)_+$, producing $h \in \mathrm{sh}^\circ(\pi_n)$). However, by lemma 6.10, the

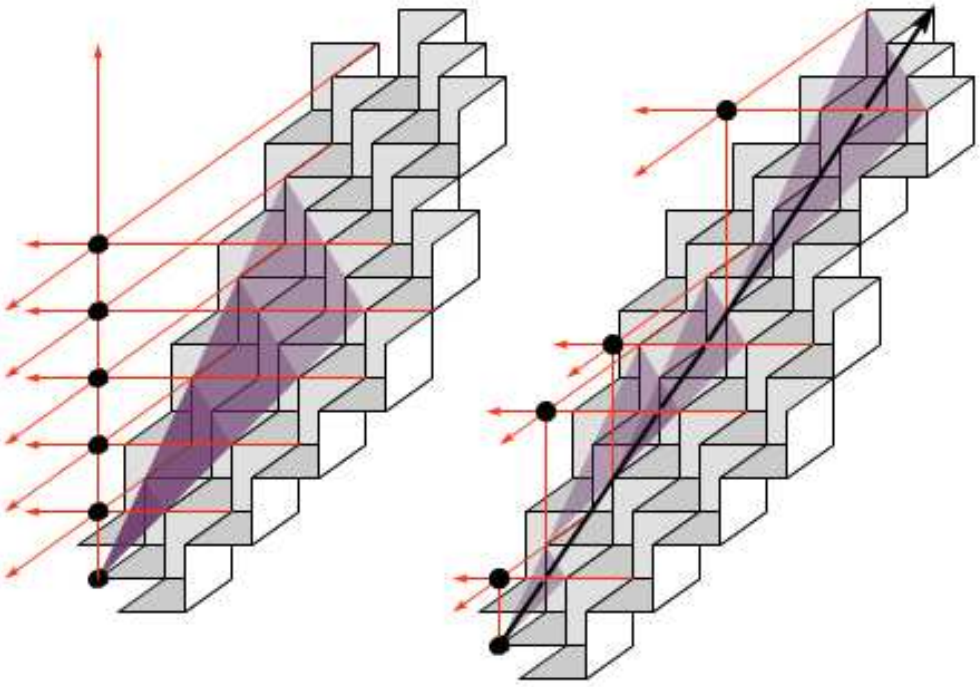

Figure 4: The hexagonal packing again. Shadows of ultrafilters receding to an unbounded distance from Π_0 (left) will increase without moving away from some common point. Shadows of ultrafilters converging (tamely) to $\rho(\xi)$, $\xi \in \partial_\infty X$ (right) will 'move faster than they grow'.

set $\mathrm{sh}^\circ(\pi_{n+1})$ is contained in $\mathrm{sh}^\circ(\pi_n)$, implying $h \notin \mathrm{sh}^\circ(\pi_{n+1})$. Applying induction we see that h then belongs to π_{n+k} for all k.
As a result, for all n we must have $\pi_n \smallsetminus \mathrm{sh}^\circ(\pi_n) \subseteq T(\xi)$. In particular, $\min(\pi_n)_-$ is contained in $T(\xi)$ for all n, as a result of applying corollary 6.10.
Let $\pi = \pi_1$, and recall that $\min(\pi)_-$ is inconsistent (lemma 6.7). This contradicts the preceding lemma when applied with $A = \min(\pi)_-$. ■

6.2.2 Finishing the proof.

Proof of theorem 6.1: Suppose $\Delta(-, \Pi_0)$ is unbounded, and let us construct a geodesic ray $(\pi_n)_{n=1}^\infty$ such that $\Delta(\pi_{n+1}, \Pi_0) > \Delta(\pi_n, \Pi_0)$ for all n. Applying proposition 6.12 will produce a point of $\Re\mathcal{H}$ outside the image of ρ. Our construction of the desired ray in Π is a pigeonhole argument slightly generalizing an argument Williams had used for the special case of Coxeter groups [Wil98].
Since $\Delta(-, \Pi_0)$ is unbounded, for every $n \in \mathbb{N}$ we select an ultrafilter π_n at a distance n from Π_0. For each π_n, let $\pi_{0,n}$ be an element of Π_0 at a distance n from π_n. Since G acts co-finitely on Π_0 and preserves $\Delta(-, \Pi_0)$, we may select a subsequence $(\pi_{n_k})_{k=1}^\infty$ such that π_{0,n_k} is the same for all k. Denote this ultrafilter by π_0.

Since $\Delta(\pi_{n_k}, \pi_0) = \Delta(\pi_{n_k}, \Pi_0) = n_k$, for each k there is a vertex-path p_k in Π of the form

$$p_k = (\pi_0, \pi_{1,n_k}, \ldots, \pi_{n_k-1,n_k}, \pi_{n_k}) \tag{85}$$

satisfying $\Delta(\pi_{i,n_k}, \Pi_0) = i$ for all relevant i.
By the local finiteness of $C(\mathcal{H})$, one may pass to a subsequence yet again, so that every path p_{k_l} is an initial subpath of the path $p_{k_{l+1}}$. The union of all the p_{k_l} is then an escaping geodesic ray in Π. ■

References

[BH93] Brigitte Brink and Robert B. Howlett. A finiteness property and an automatic structure for coxeter groups. *Mathematische Annalen*, (296):179–190, 1993.

[BH99] M. R. Bridson and A. Haefliger. *Metric Spaces of Non-Positive Curvature*, volume 319 of *Grundlehren der mathematischen Wissenschaften*. Springer Verlag, Berlin-Heidelberg, 1999.

[Bow98] B.H. Bowditch. Cut points and canonical splittings of hyperbolic groups. *Acta Math.*, 180:145–186, 1998.

[Cap05] P. E. Caprace. Conjugacy of one-ended subgroups of coxeter groups and parallel walls (preprint). *http://www.arXiv.org:math/0508057*, 2005.

[CK00] S. Croke and B. Kleiner. Spaces with non-positive curvature and their ideal boundaries. *Topology*, 39:549–556, 2000.

[CNar] Indira Chatterji and Graham A. Niblo. From wall spaces to CAT(0) cube complexes. *International Journal of Algebra and Computation*, to appear.

[DS91] M.W. Davis and M.D. Shapiro. Coxeter groups are automatic. *preprint MRI 91-15*, 1991.

[Gur05] Dan Guralnik. Ends of cusp-uniform groups of locally connected continua – I. *International Journal of Algebra and Computation*, 15(4):765–798, August 2005.

[Gur06] Dan Guralnik. Local finiteness of cubulations for cat(0) groups. *preprint (submitted to Geometry, Groups and Dynamics) http://www.arXiv.org/abs/math.GR/0610950*, 2006.

[MR99] Michael Mihalik and Kim Ruane. CAT(0) groups with non-locally connected boundary. *Journal of the London Mathematical Society*, 60(3):757–770, 1999.

[MR01] Michael Mihalik and Kim Ruane. CAT(0) HNN-extensions with non-locally connected boundary. *Topology and its Applications*, 110(1):83–98, 2001.

[Nic04] Bogdan Nica. Cubulating spaces with walls. *Algebraic and Geometric Topology*, 4:297–309, May 2004.

[NR98] Graham H. Niblo and Lawrence D. Reeves. The geometry of cube complexes and the complexity of their fundamental groups. *Topology*, 37(3):621–633, 1998.

[NR03] Graham H. Niblo and Lawrence D. Reeves. Coxeter groups act on cat(0) cube complexes. *Journal of Group Theory*, (6):309–413, 2003.

[Ont05] P. Ontaneda. Cocompact cat(0) spaces are almost geodesically complete. *Topology*, 44:47–62, 2005.

[Rol98] M.A. Roller. Poc sets, median algebras and group actions. *University of Southampton, Faculty of Math. stud., preprint series*, 1998.

[Sag95] Michah Sageev. Ends of groups pairs and non-positively curved cube complexes. *Proc. London Math. Soc.*, 3(71):586–617, 1995.

[Wil98] B.T. Williams. *Two topics in geometric group theory.* PhD thesis, University of Southampton, 1998.

[Wis04] Daniel T. Wise. Cubulating small-cancellation groups. *GAFA*, 14(1):150–214, 2004.